Circuit Technologies for Multi-Core Processor Design

Stefan Rusu

Intel Corporation

Santa Clara, CA

stefan.rusu@intel.com

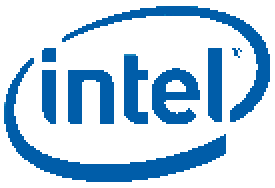

Outline

- Dual-core architectural directions
- Interconnect trends
- Power and leakage reduction
  - Cache Sleep and Shut-off Modes
  - Long-Le Transistor Usage
- Voltage domains
- Clock distribution
- Package details
- DFT/DFM features
- Thermal management
- Summary

Dual Core Processors Everywhere!

Intel tips Viiv, Yonah in consumer push

Mark LaPedus, 01/06/2006

SAN JOSE, Calif. — Making a big splash in the consumer market, Intel Corp. on Thursday unveiled a dual-core processor, PC platform and several content alliances that are said to provide the foundation for digital entertainment and wireless laptops.

Laptop

Intel unwraps dual-core Xeon server processors

Tom Krazit, 10/10/2005

At an event Monday in San Francisco, Intel unveiled its first dual-core Xeon chips for two-processor and four-processor servers, previously known by the Paxville code name.

Server

Intel ready to ship dual-core processors

Daniel A. Begun, 3/15/2005

Earlier this year, Intel announced that desktop processors using dual-core technology will be available by the end of June, but the company recently hinted at March's Intel Developer Forum (IDF) that dual-core processors could be available even sooner than that.

Desktop

Tulsa – Xeon® Processor

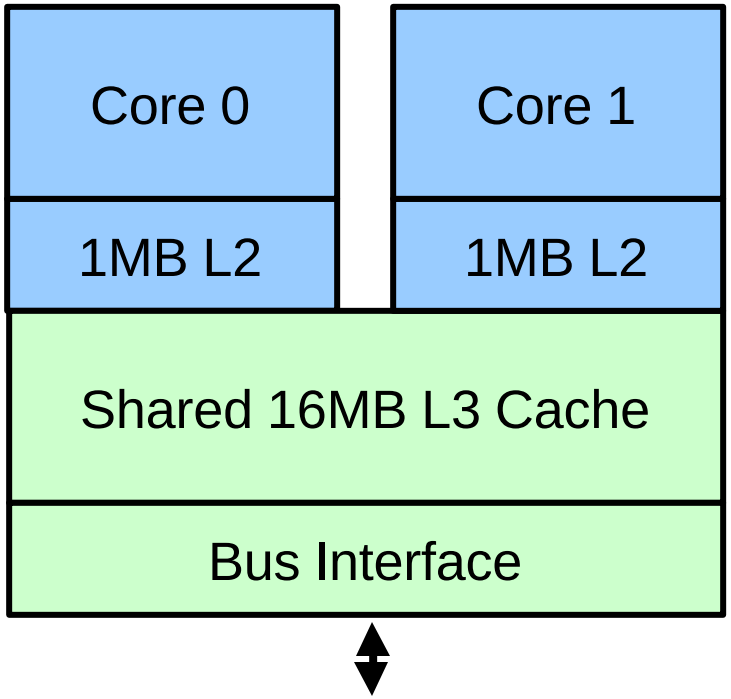

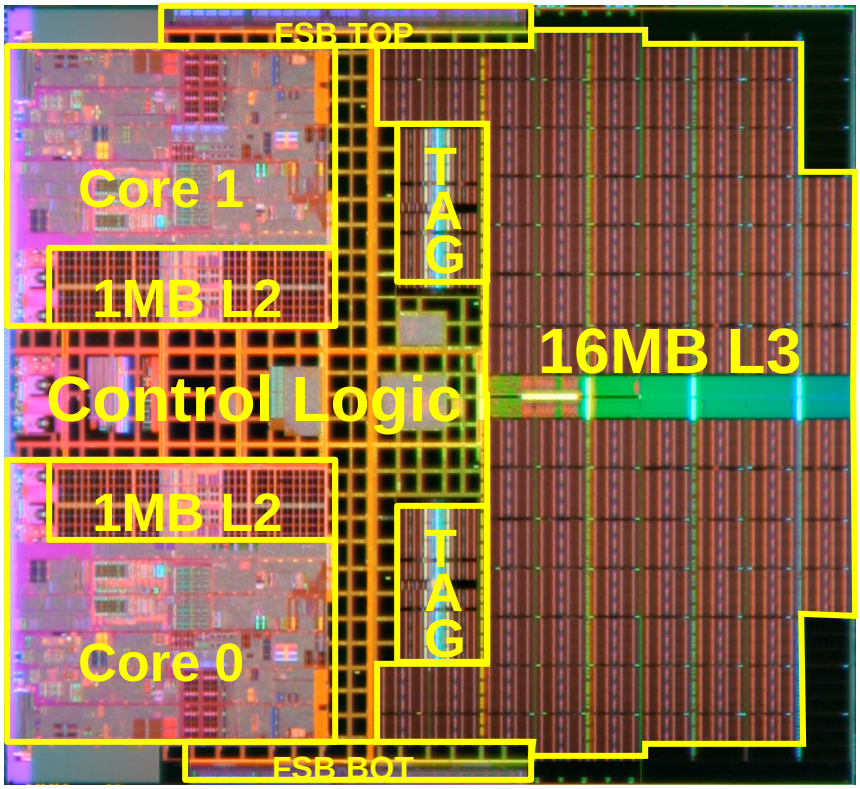

- Process technology: 65 nm, 8 Cu interconnect layers
- Transistor count: 1.328 Billion
- Die area: 435 mm^2

Montecito – Itanium® 2 Processor

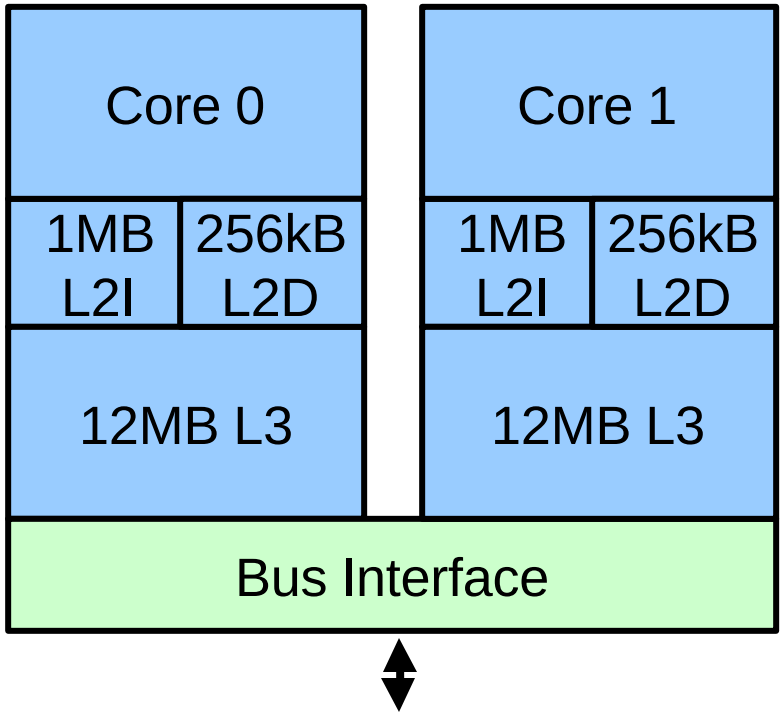

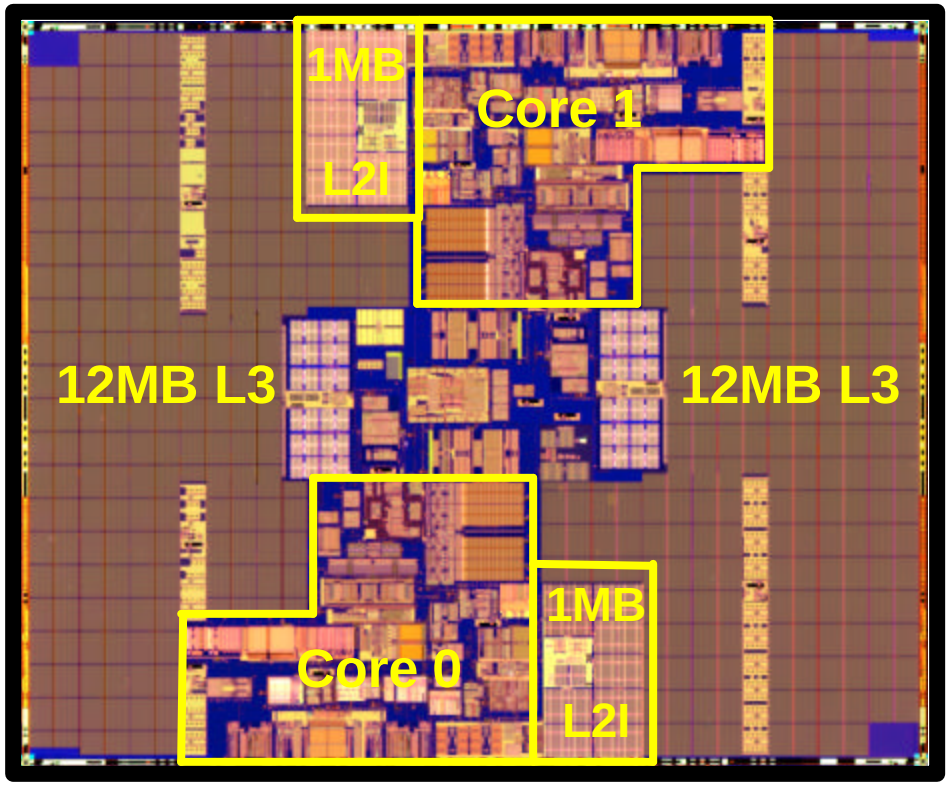

- Process technology: 90nm, 7 Cu interconnect layers
- Transistor count: 1.72 Billion
- Die area: 596mm^2

Yonah – Mobile/Desktop/Blade Server

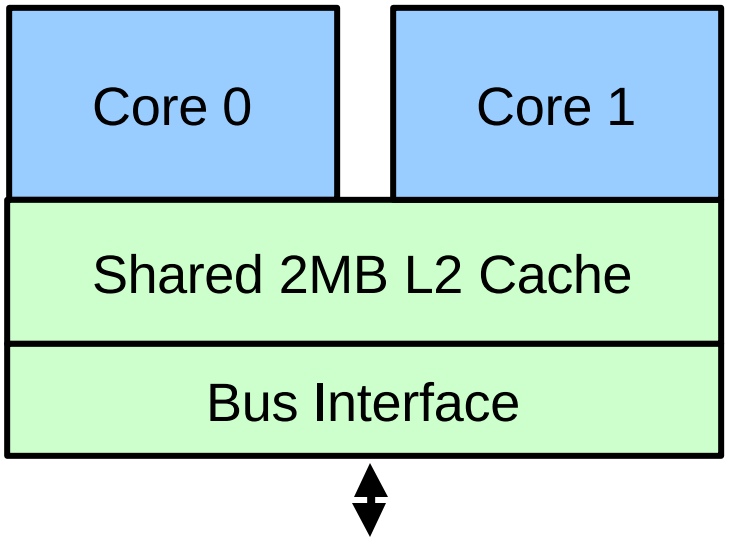

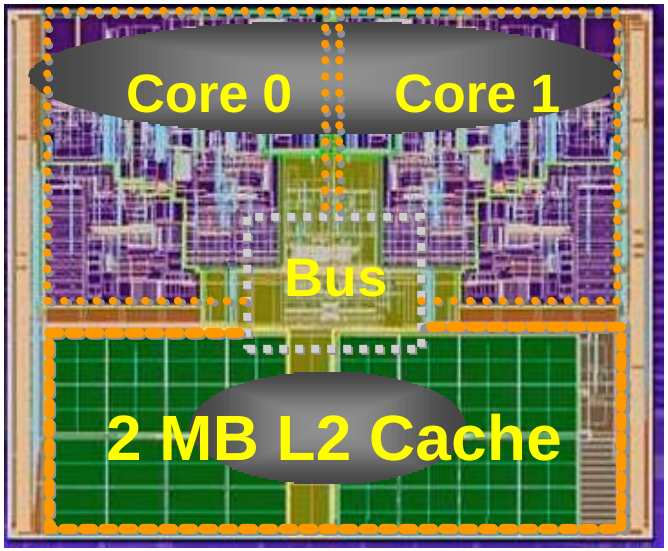

- Process technology: 65 nm, 8 Cu interconnect layers
- Transistor count: 151 Million
- Die area: 90.3 mm^2

Moore's Law for Multi-Core Processors

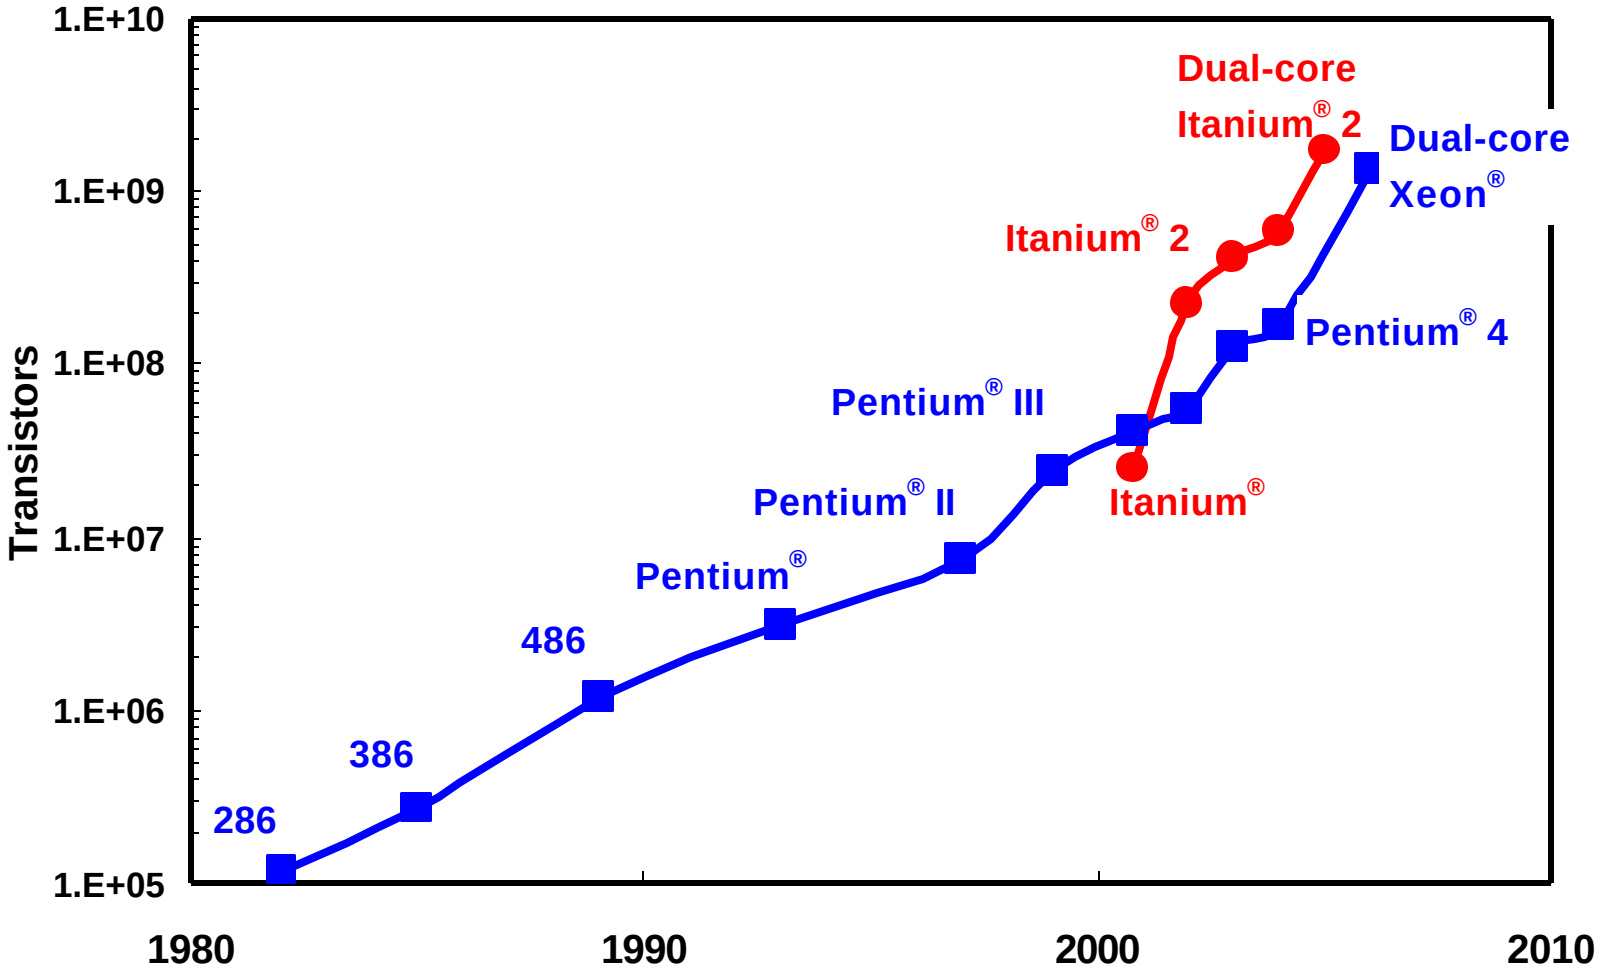

Process technology scaling enables the number of cores to double every two years

Server Processors L3 Cache Trend

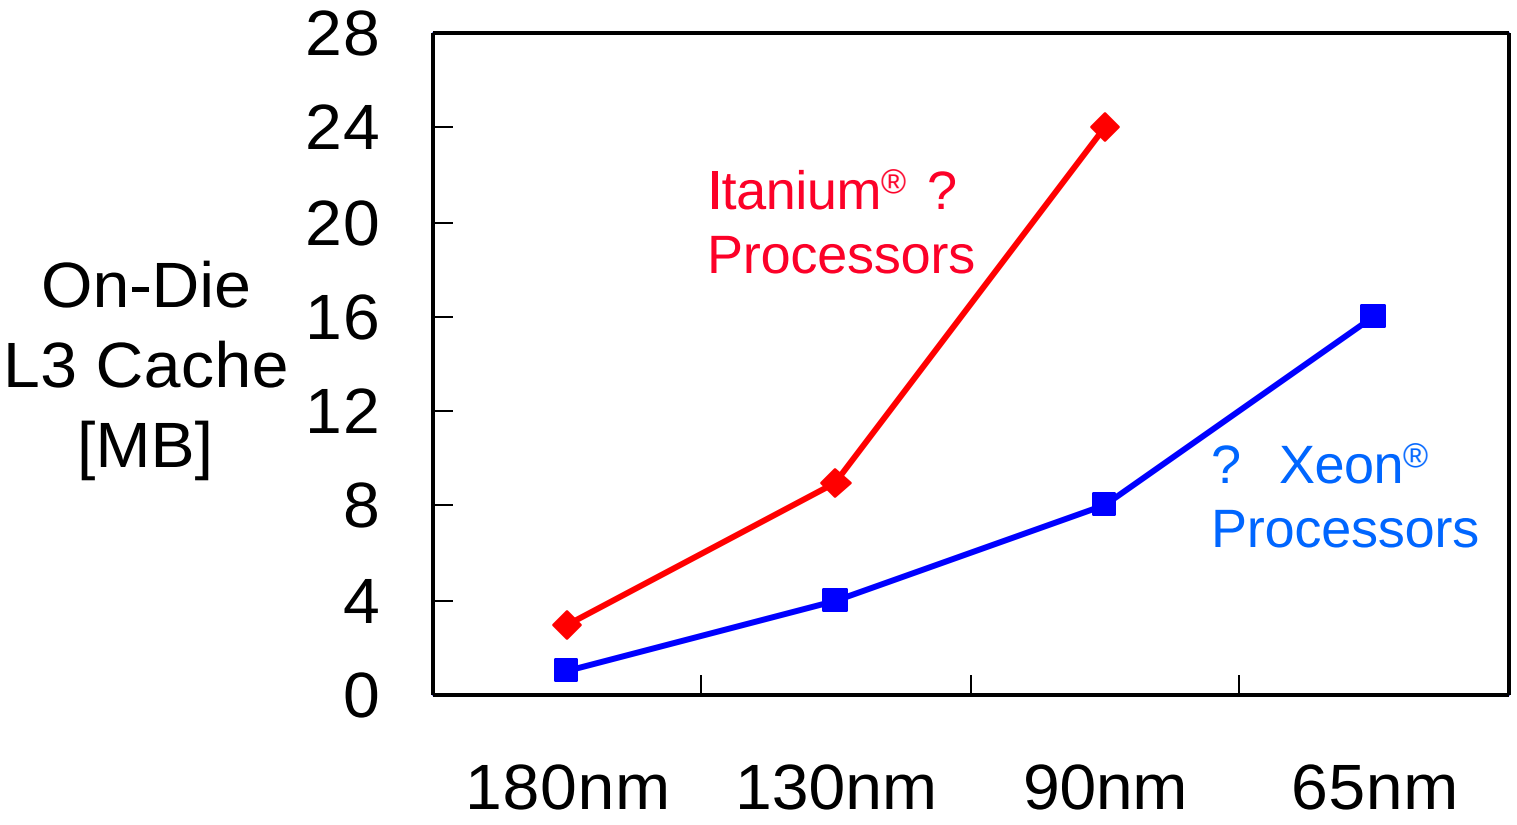

Cache size doubles with every process generation

SRAM Cell Size Scaling

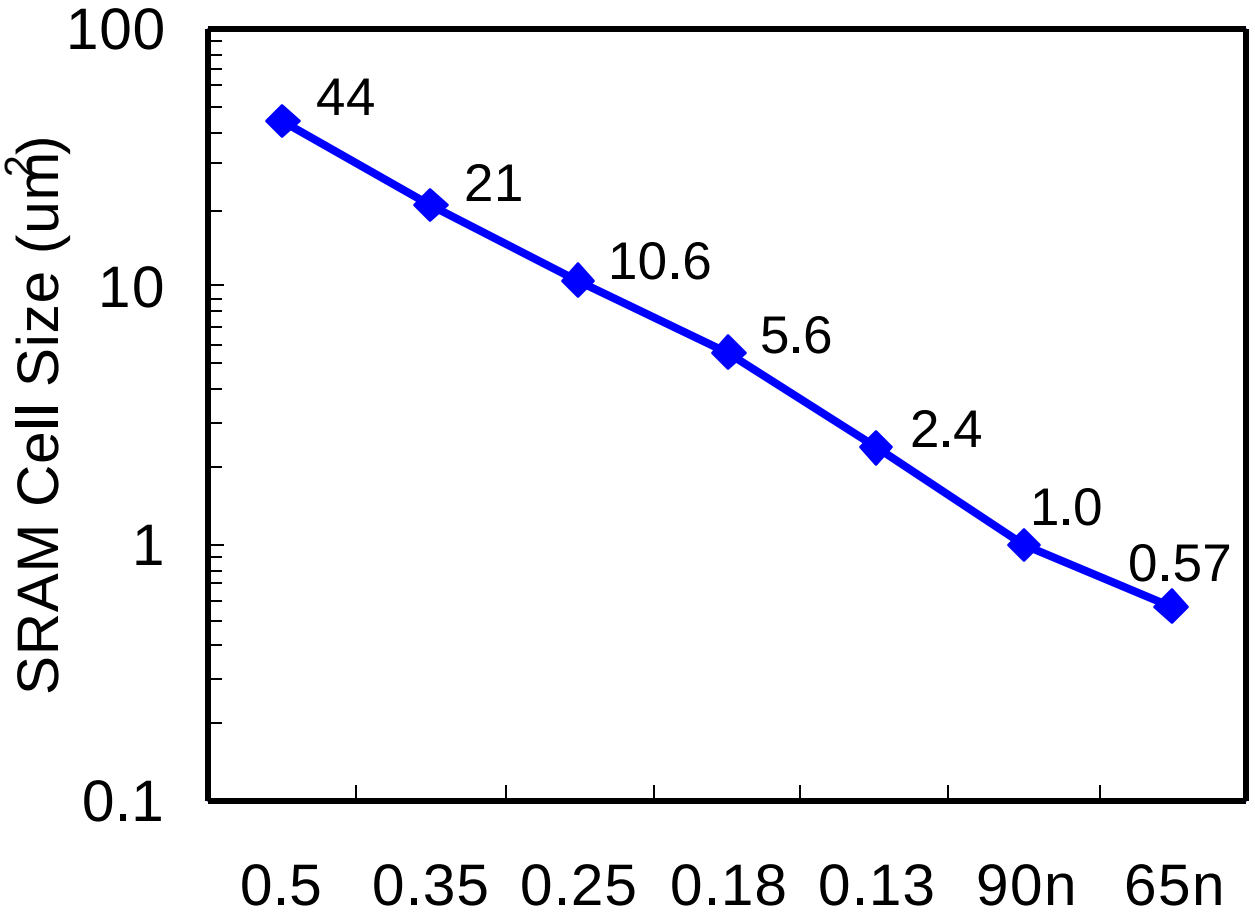

SRAM cell size scales ~0.5x per generation

Server Processors L3 Cache Trend

Dual-core Processor	Process	L3 cache per core [MB]	L3 cache per thread [MB]
Itanium®	90nm	12	6
Xeon®	65nm	8	4

Large per-core and per-thread caches enable server-class performance

On-chip Interconnect Trend

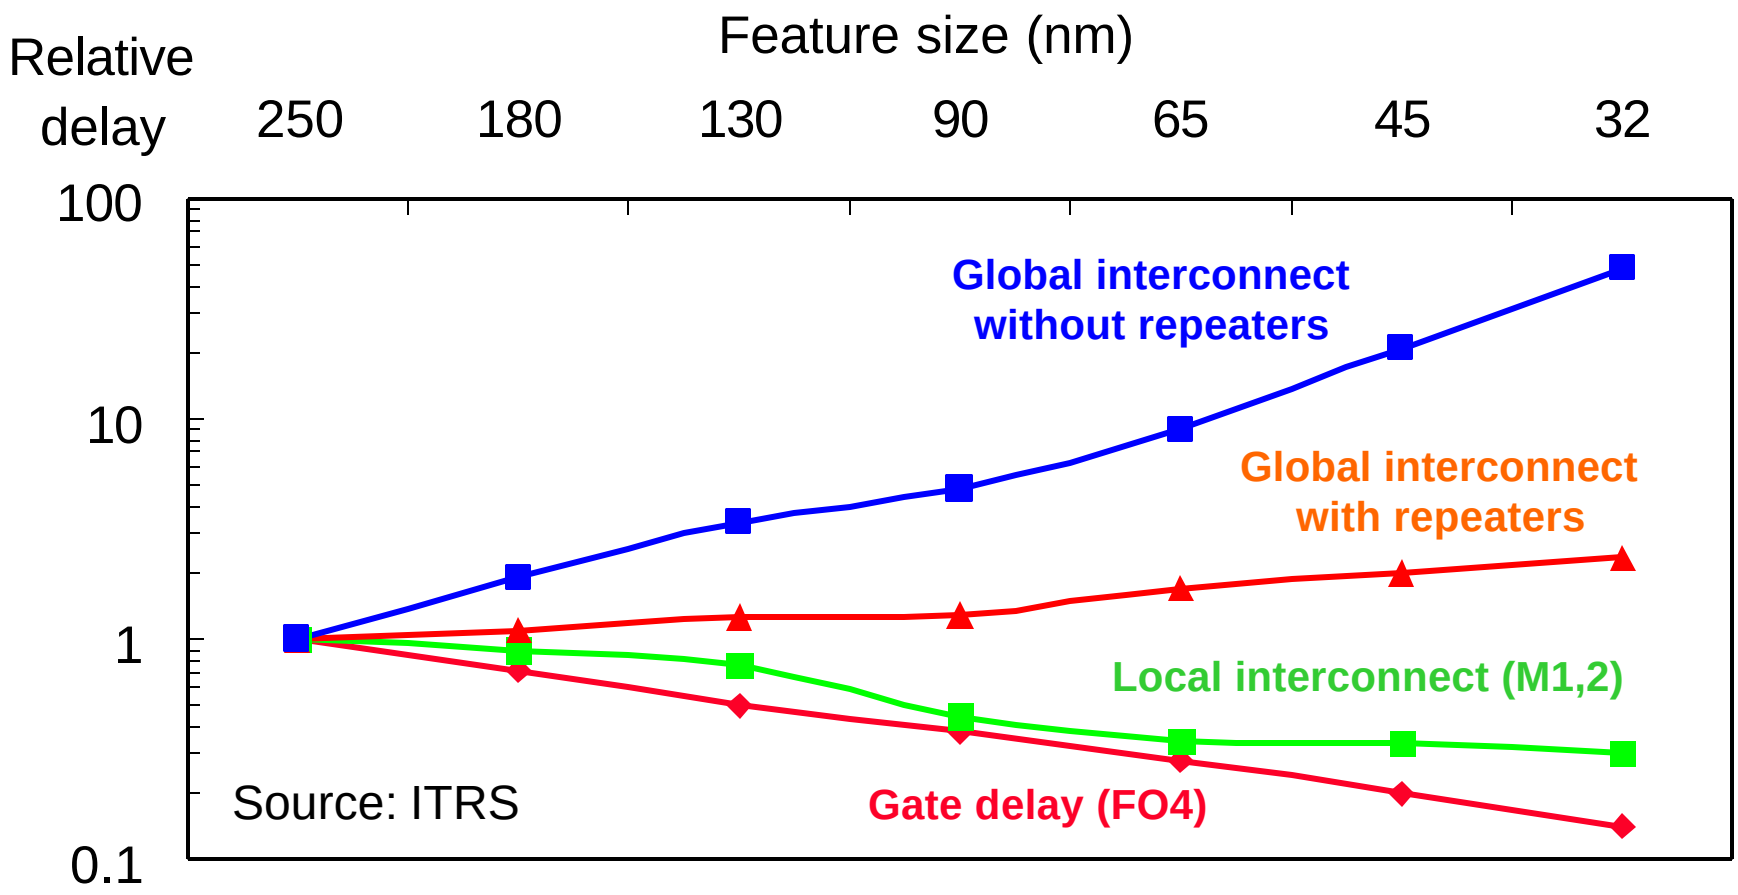

- Local interconnects scale with gate delay
- Global interconnects do not scale

Layout Techniques to Reduce Coupling

- Interleave bussed signals with other busses that switch at a different time
  - Eliminates capacitive coupling and reduces inductive noise

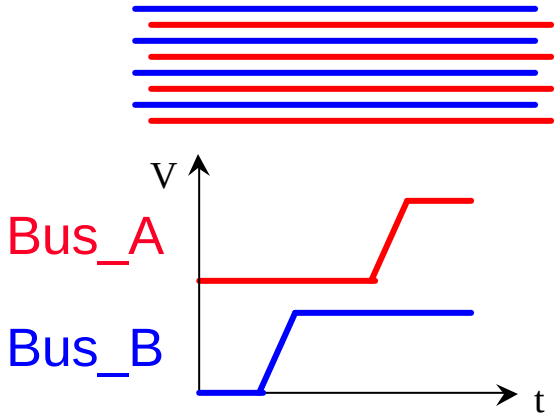

- Switch bit order of bussed signals at every turn
  - Noise will not be additive across the entire bus route

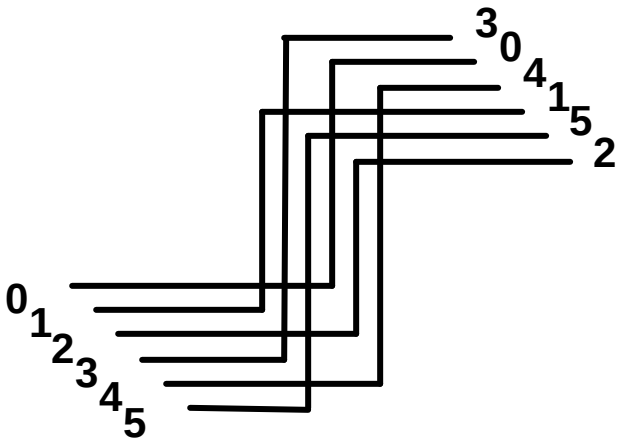

Swizzled Bus Route

Layout Techniques (Cont)

- Interleave narrow Vcc/Vss lines through bus

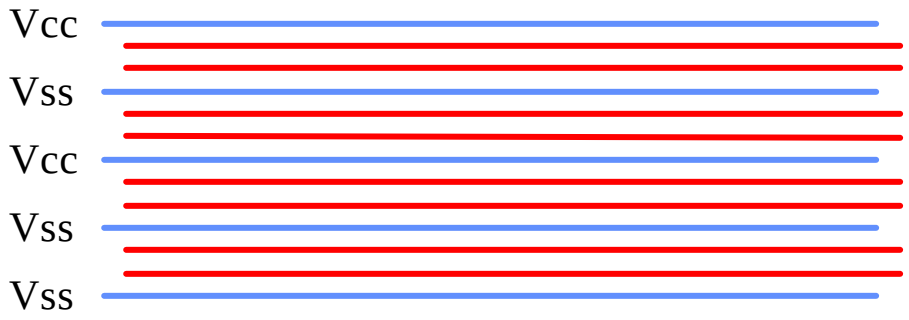

- Use staggered inverting buffers [7]

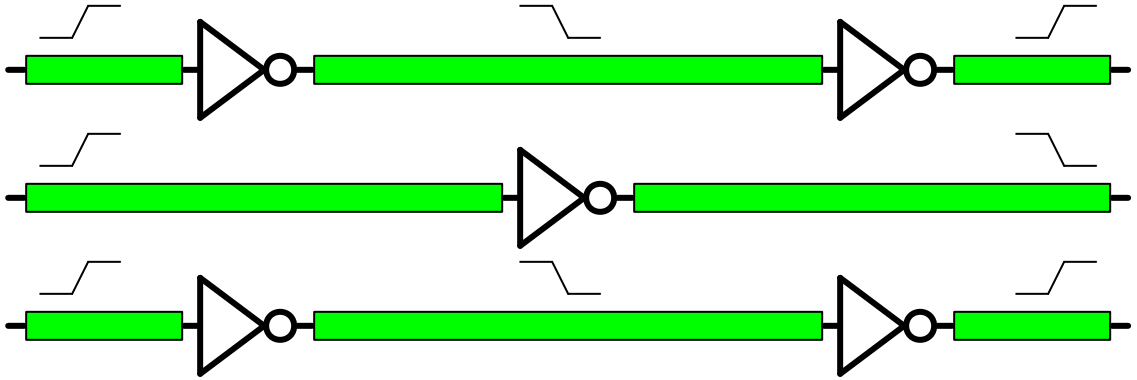

L3 Cache Sleep and Shut-off Modes

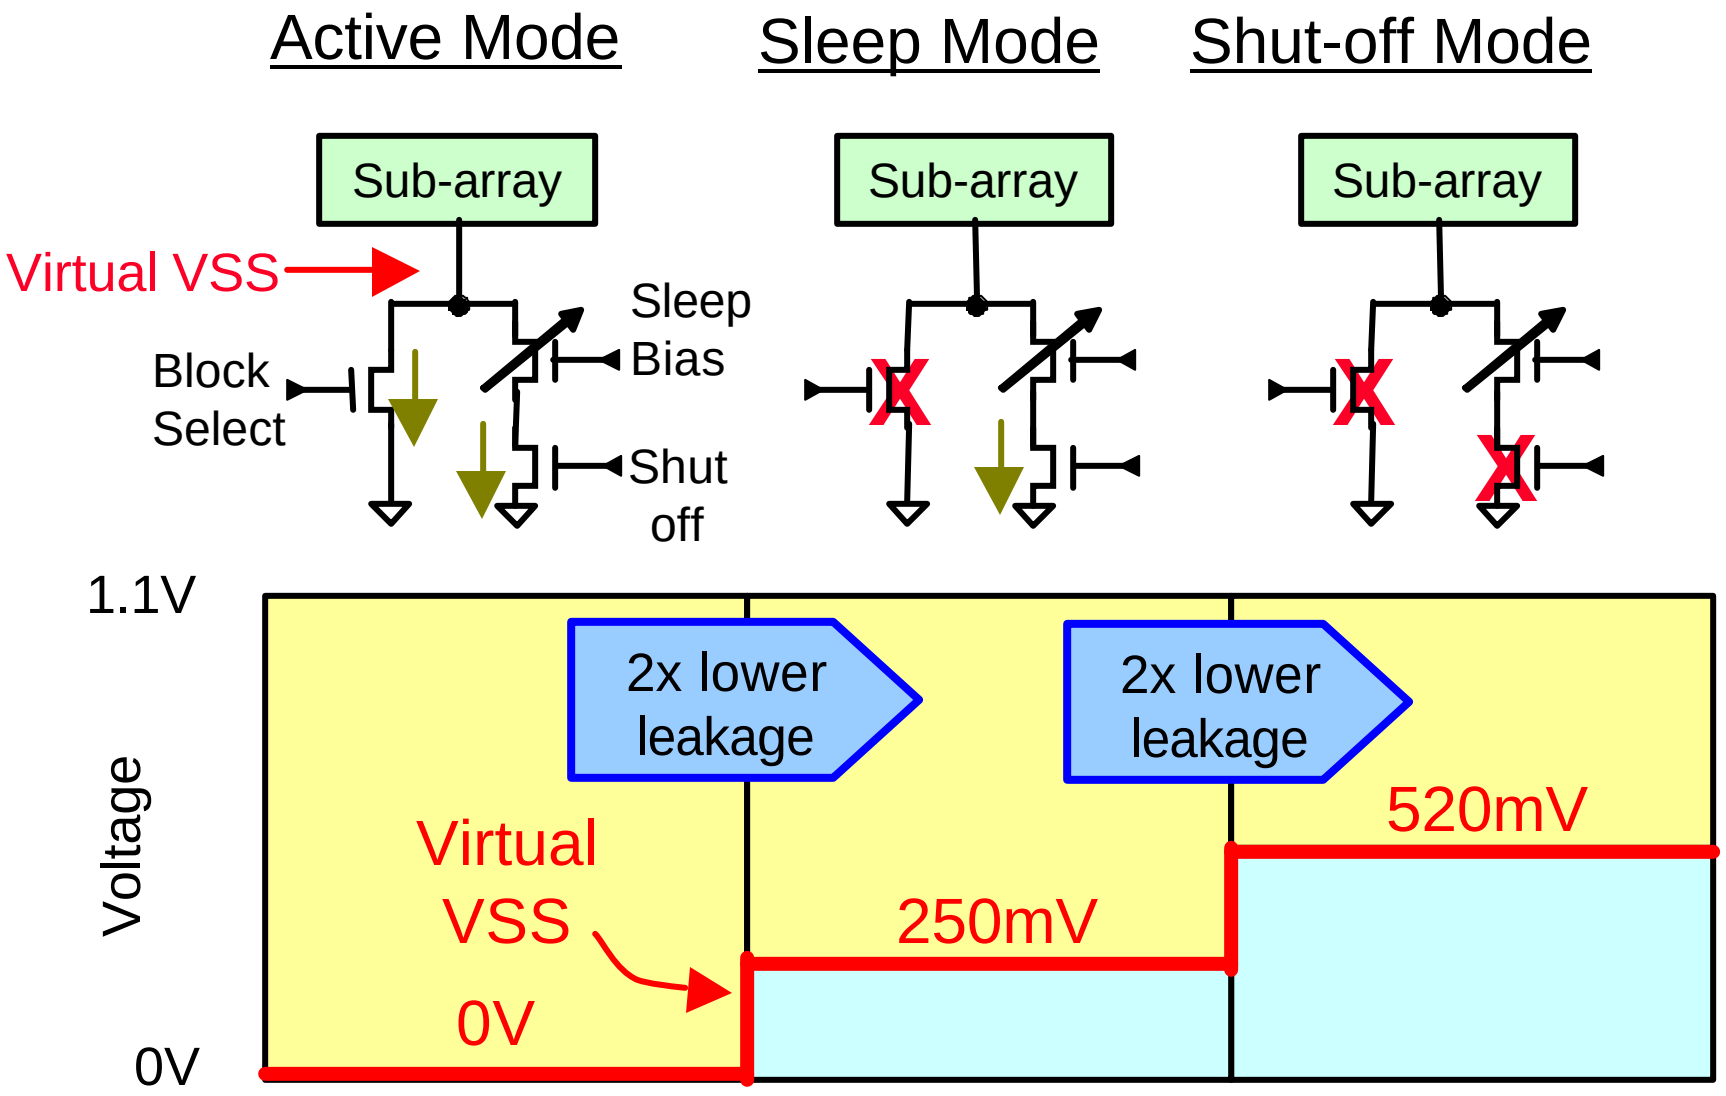

Leakage Shut-off Infrared Images

16MB SKU

All 16MB in sleep mode

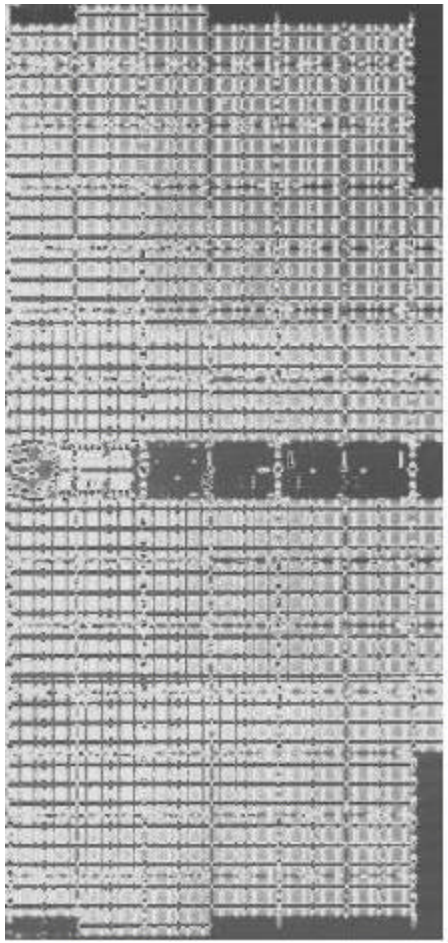

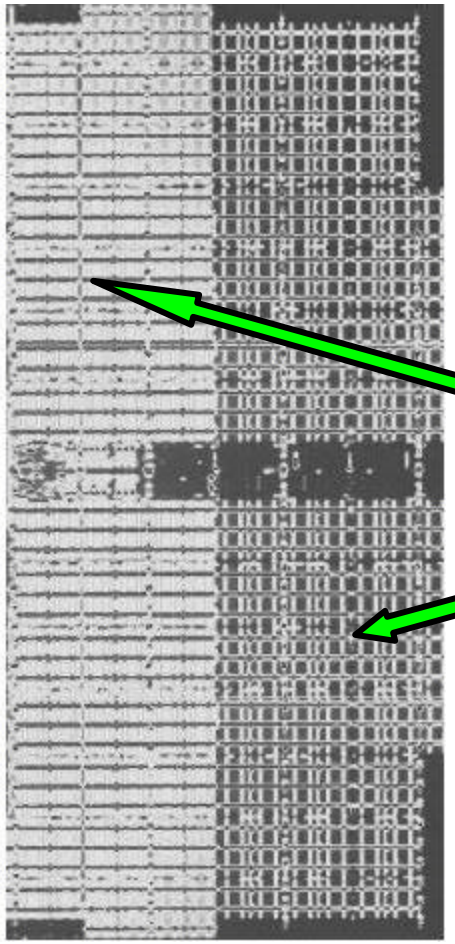

8MB SKU

8MB in sleep mode

8MB in shut-off mode

Shut-off feature reduces the leakage of the 8MB disabled sub-arrays by about 3W

Dynamic Intel® Smart Cache Sizing

- First implementation in Yonah dual-core mobile processor
  - Dynamic implementation of the shut-off mode
- HW based algorithm predicts cache usage requirements
  - Considers the % of time the CPU is in Active state compared to the various sleep states
- During periods of low activity or inactivity the processor dynamically adapts its effective cache size
  - Cache content is gradually flushed to system memory
  - Cache ways are gradually turned off (physically as well as logically), thus reducing power
- Cache ways are re-powered on demand to deliver full performance when needed

Shut-off Mode Scaling Trends

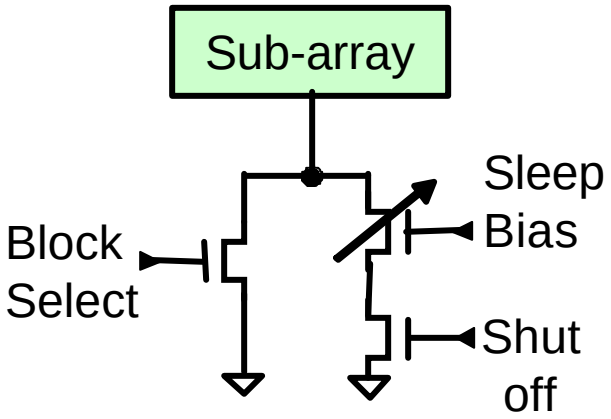

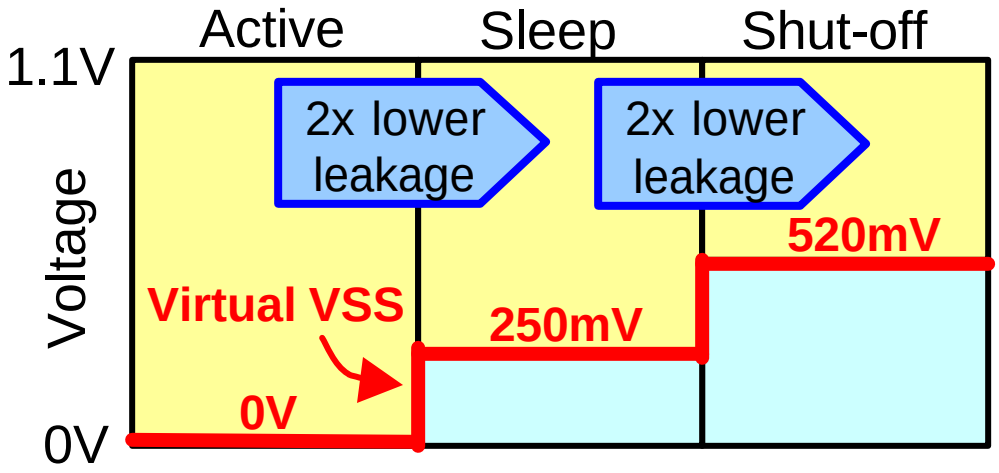

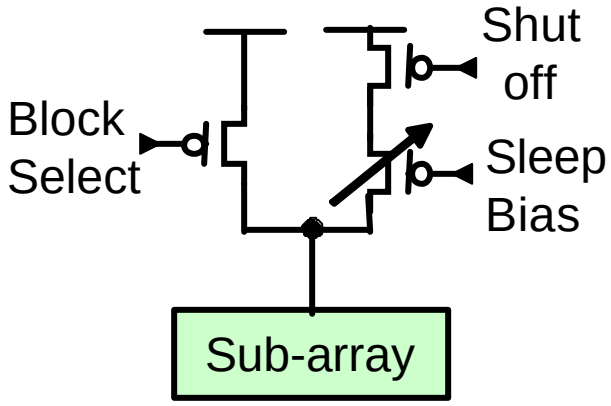

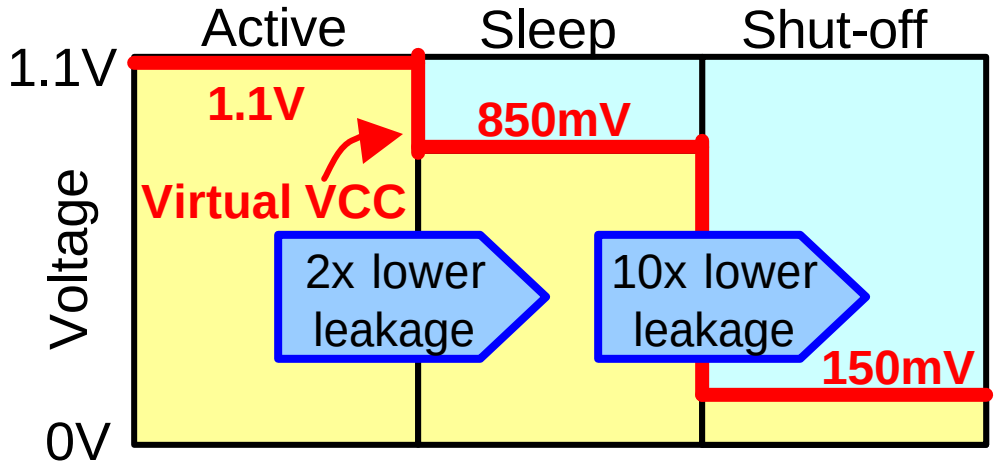

PMOS has better shut-off but larger area penalty

Leakage Mitigation: Long-Le Transistors

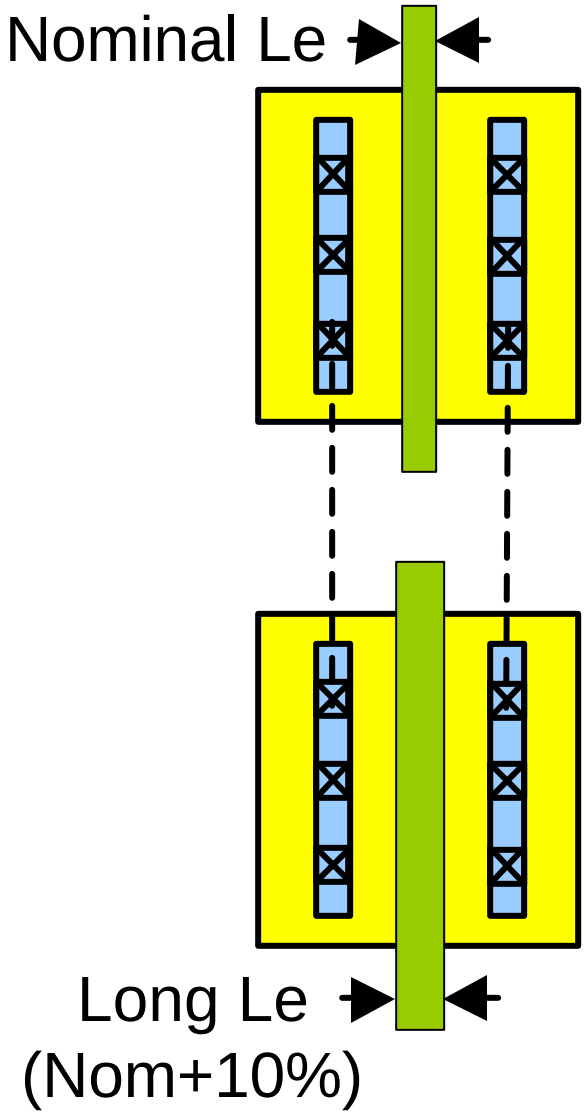

- All transistors can be either nominal or long-Le
- Most library cells are available in both flavors
- Long-Le transistors are about 10% slower, but have 3x lower leakage
- All paths with timing slack use long-Le transistors
- Initial design uses only long channel devices

Long-Le Transistors Usage Map

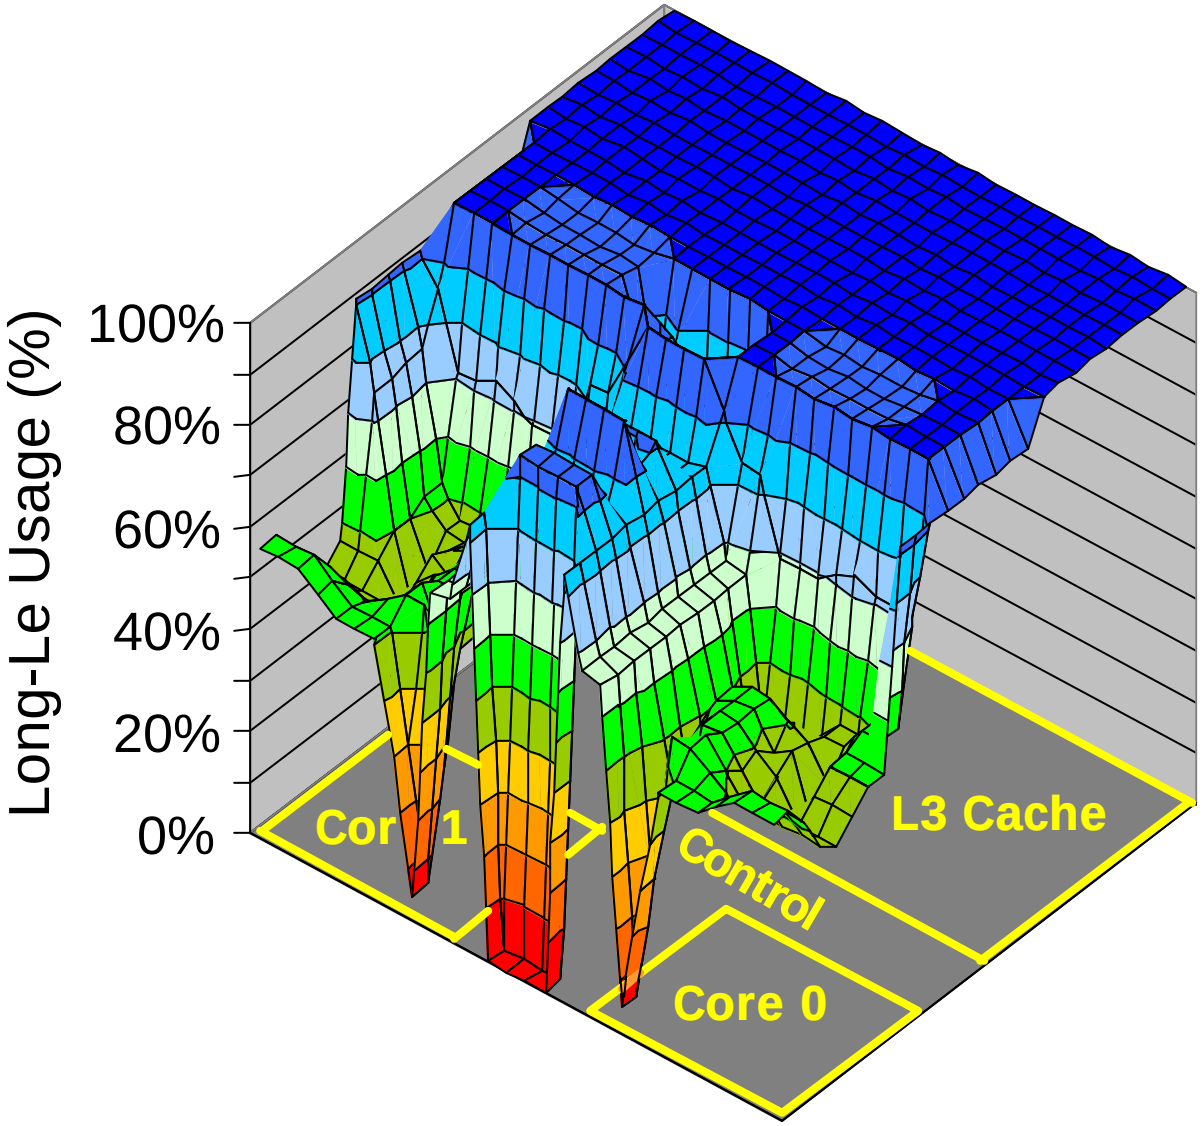

Long-Le Transistors Summary

Percentage of Long-Le device width excluding RAM arrays:

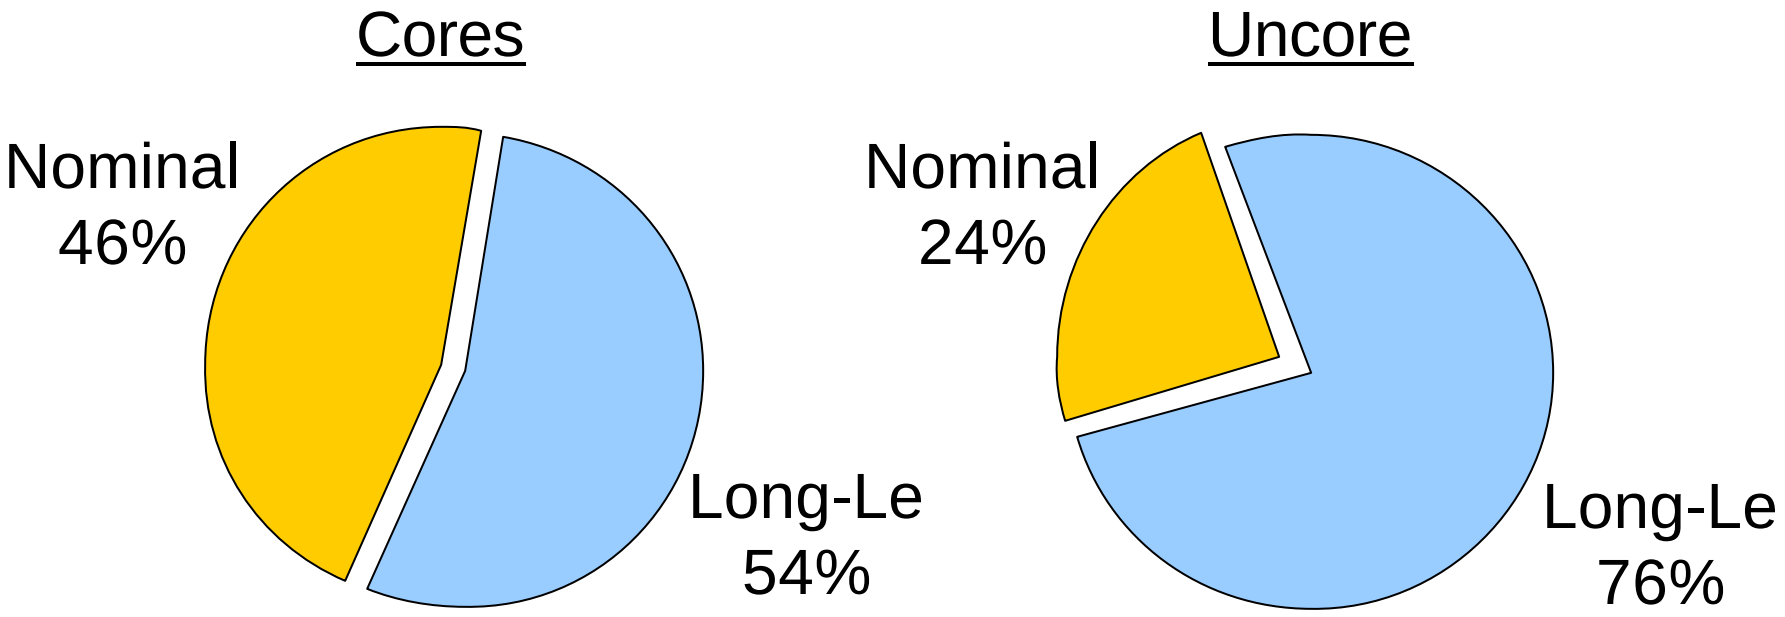

To reduce sub-threshold leakage, most devices will be slower and only a handful of transistors will be fast

Stack Forcing

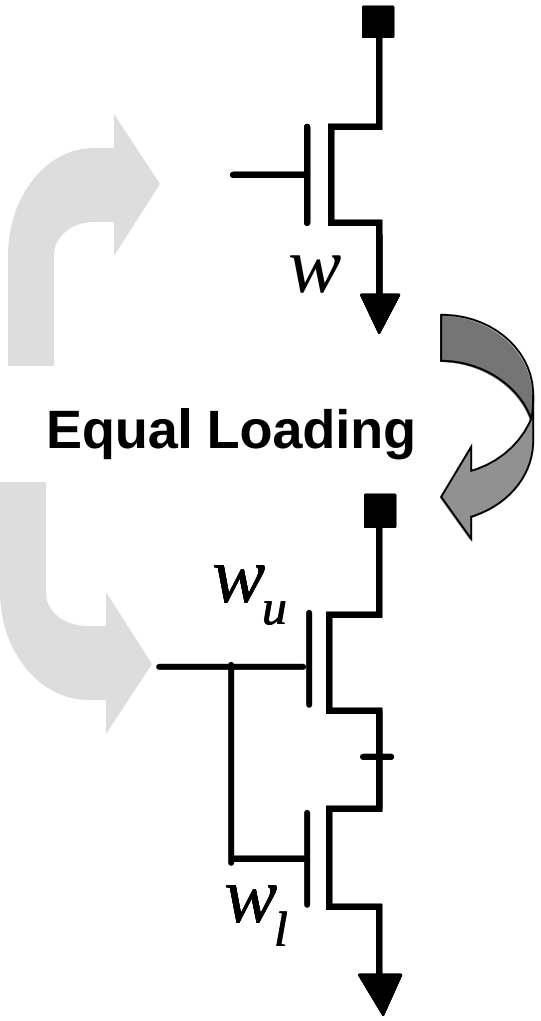

[Narendra, et al – ISLPED 2001]

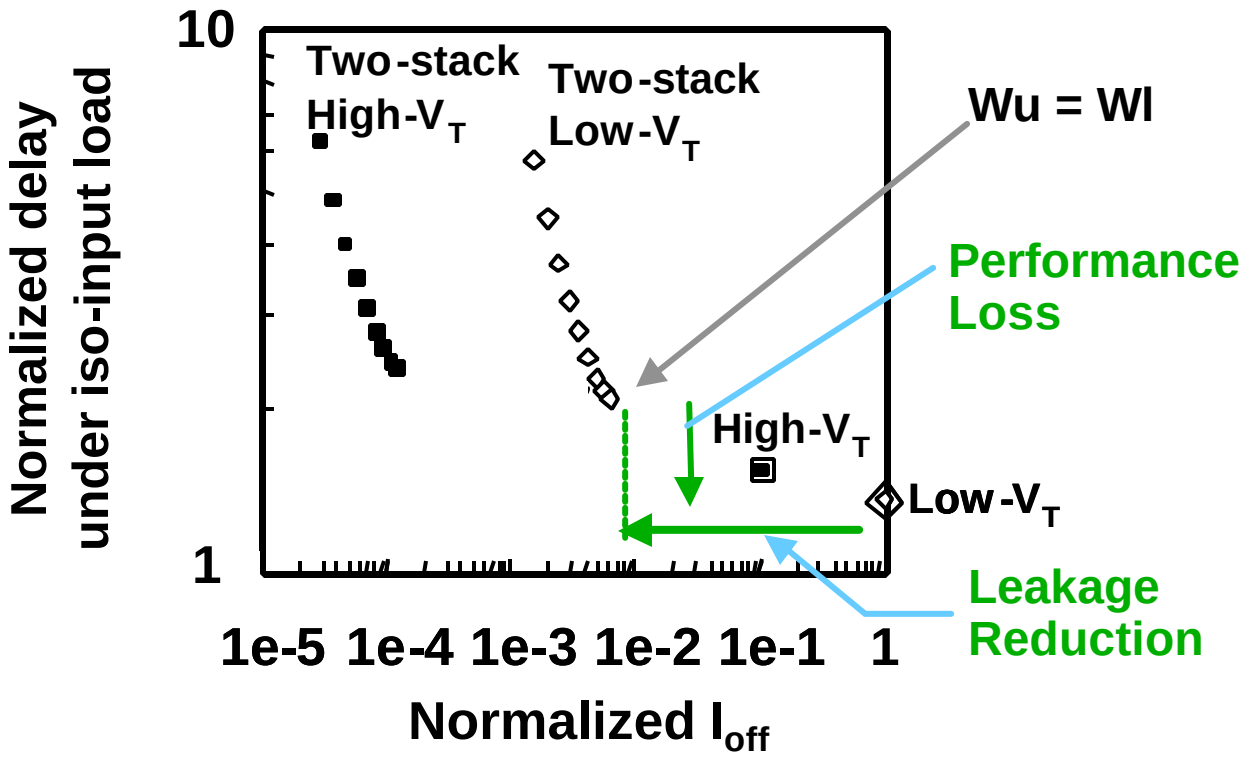

- Force one transistor into a two transistor stack with the same input load
- Can be applied to gates with timing slack
- Trade-off between transistor leakage and speed

Montecito Voltage Domains [2]

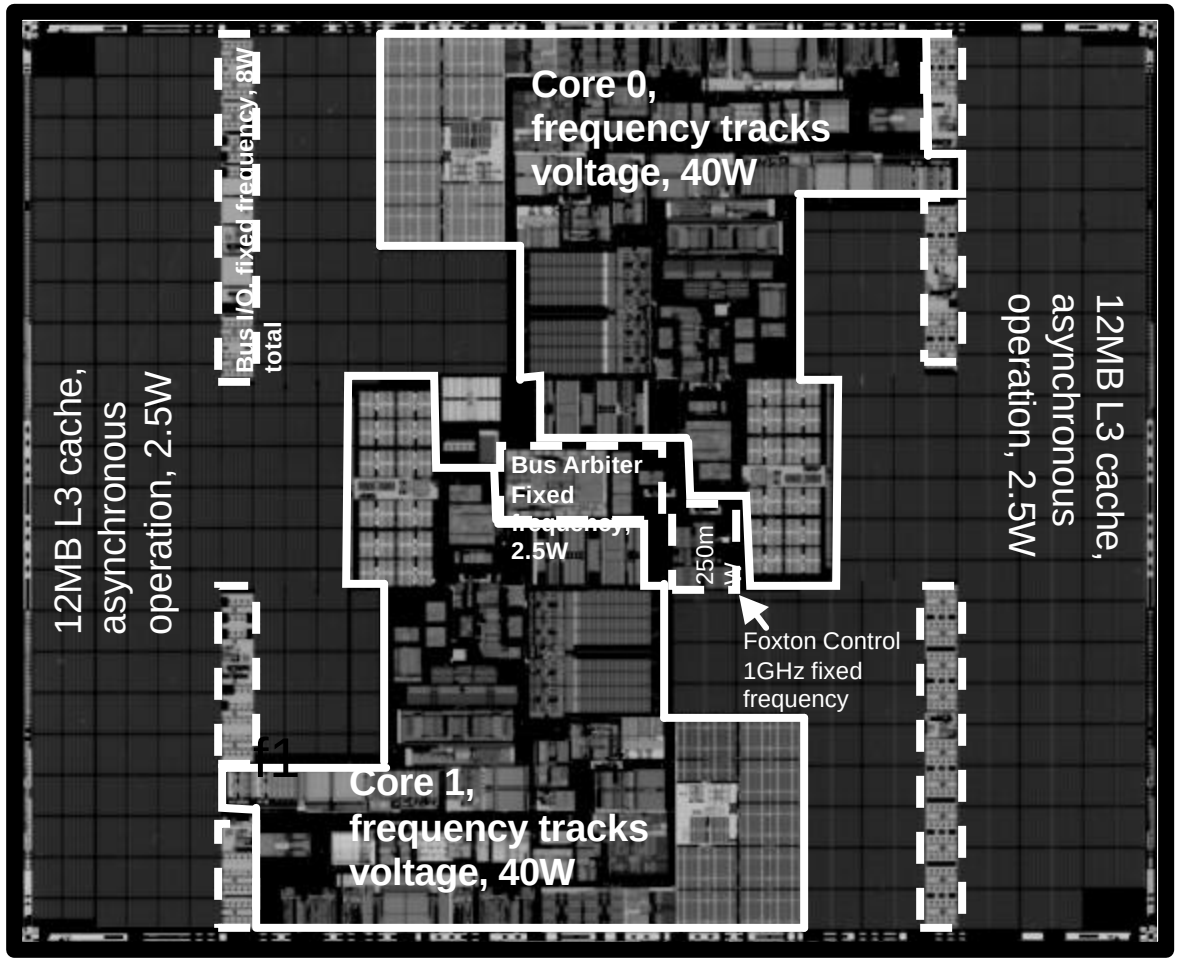

Tulsa Voltage Domains

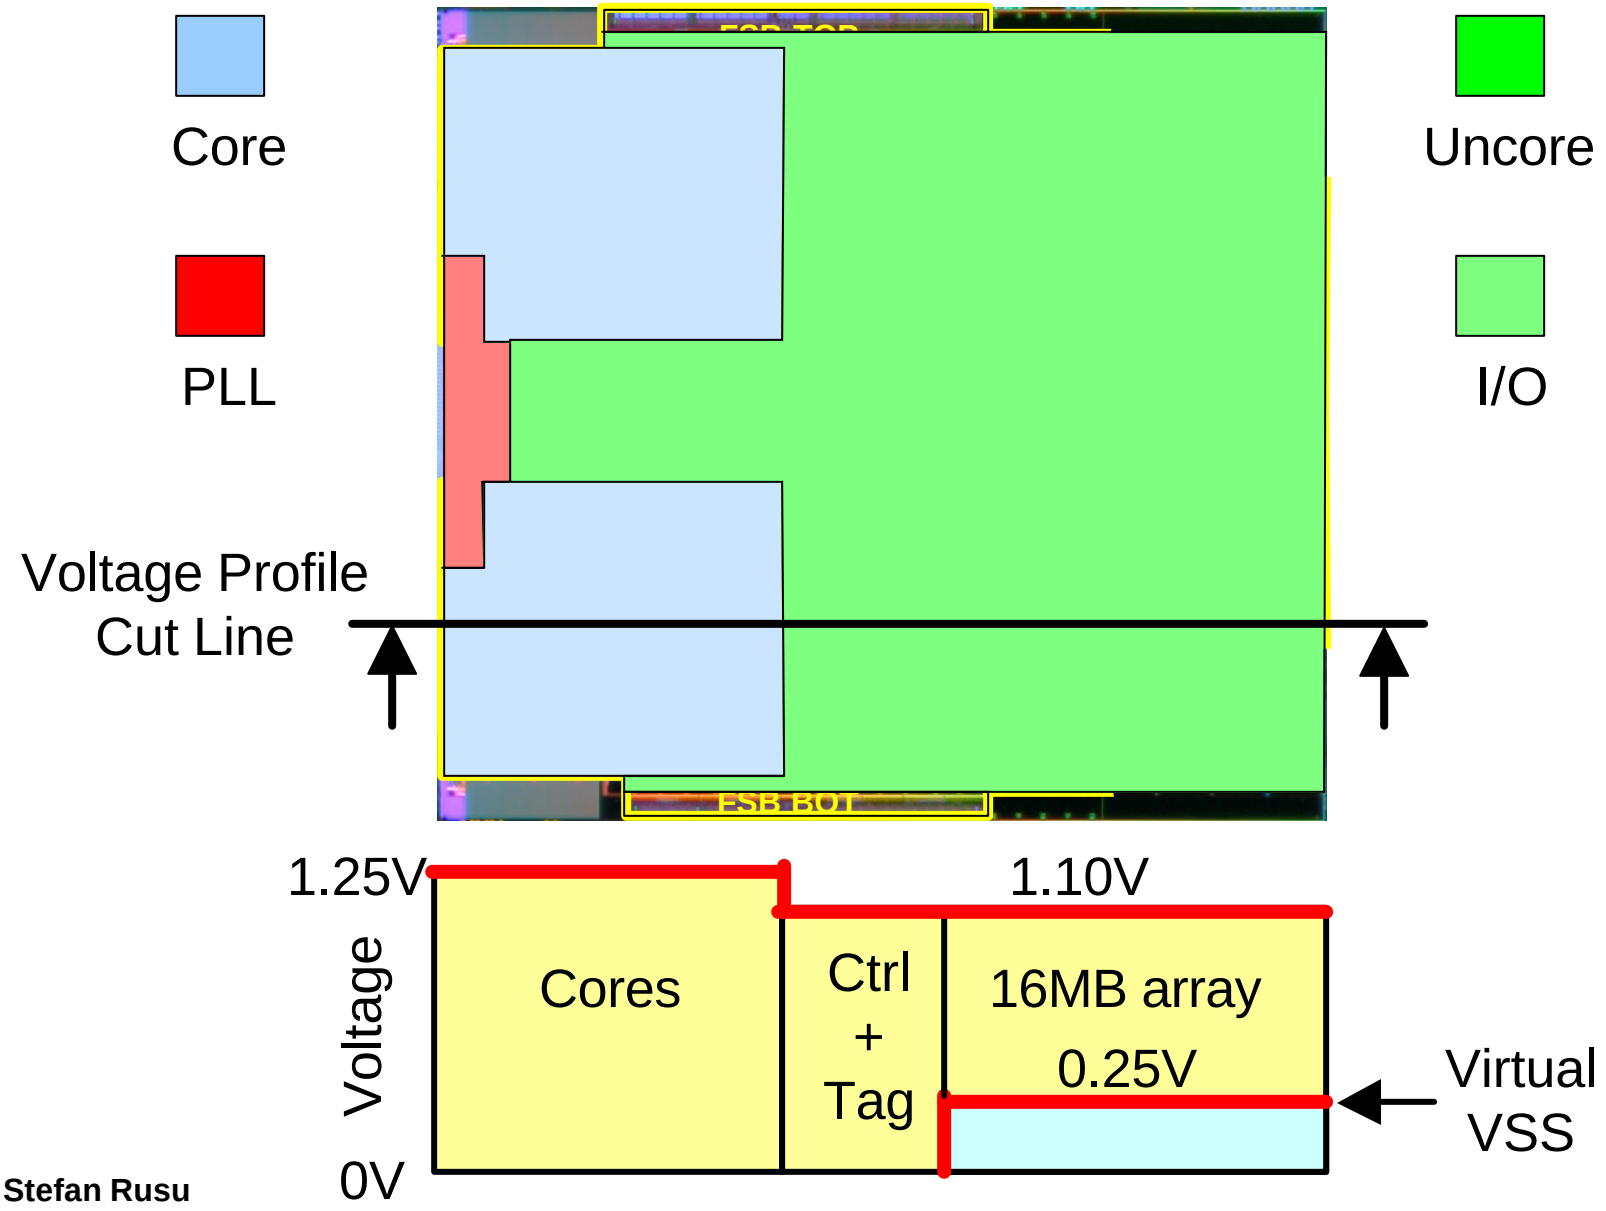

Tulsa Clock Domains [4]

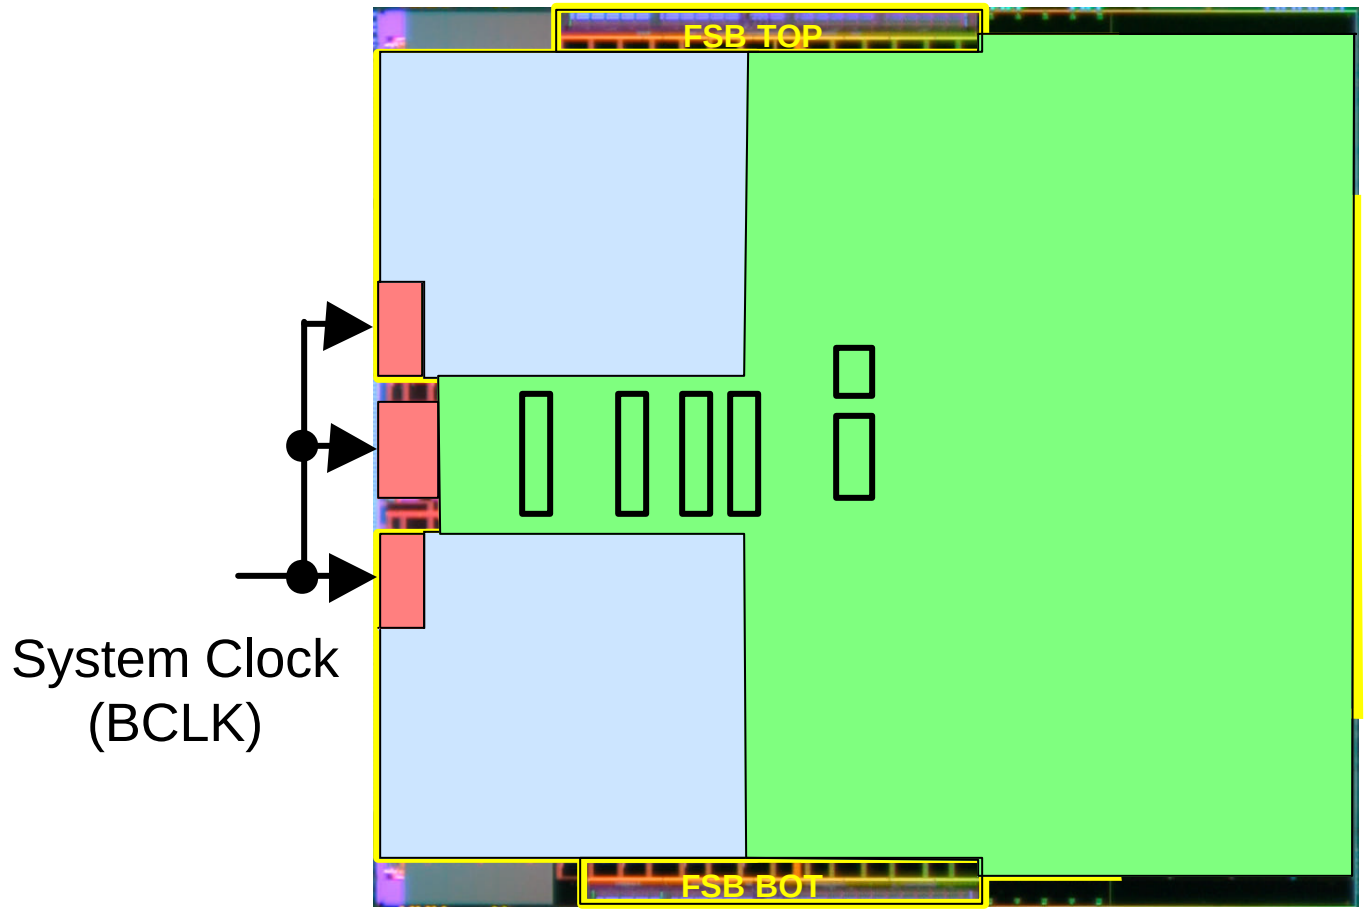

Legend:

Core PLL Uncore I/O

Tulsa Uncore Clock Distribution [4]

- **PLL (Clock Generator)**
- **Core dense MCLK grid**
- **Un-Core ZCLK grid**
- **Un-Core pre-global ZCLK spine**
- **Un-Core sparse SCLK grid**
- **Un-Core pre-global MCLK spine**
- **De-skew buffer**

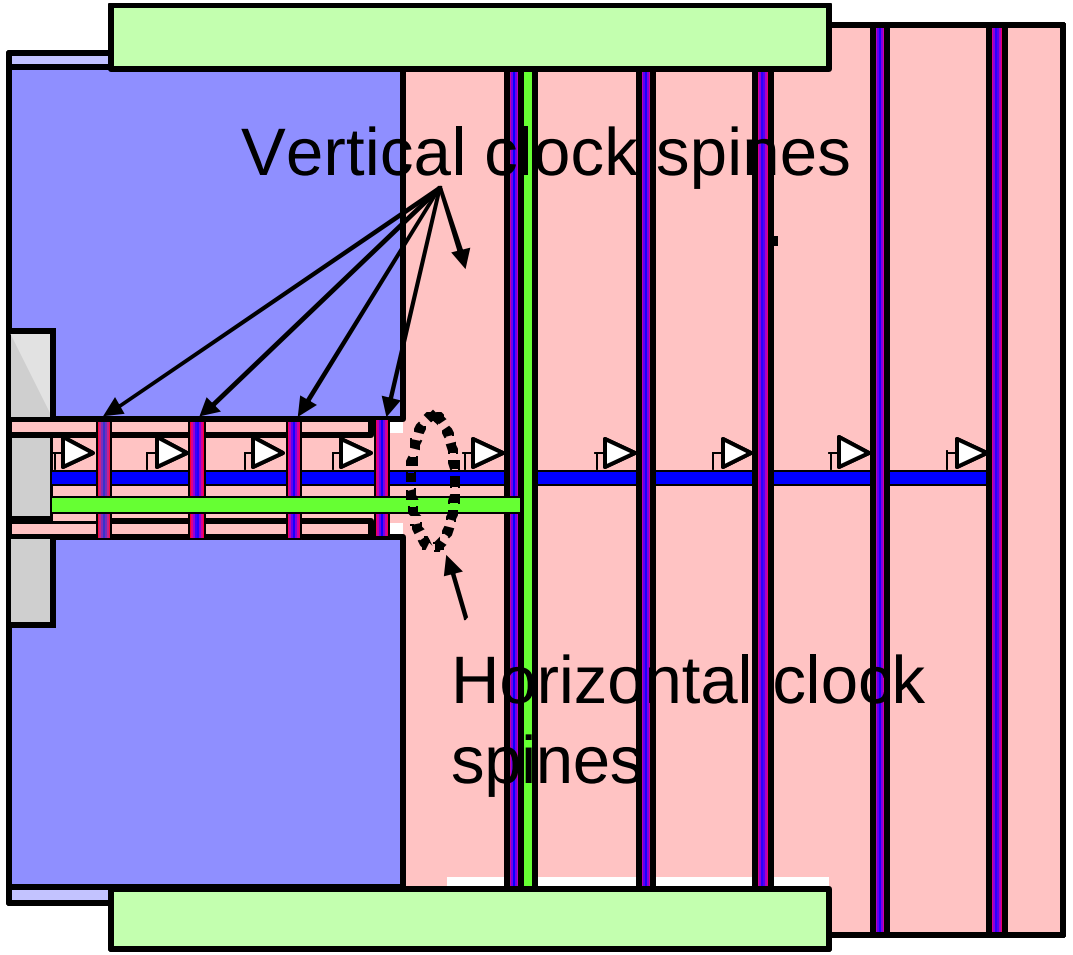

Montecito Clock System [5]

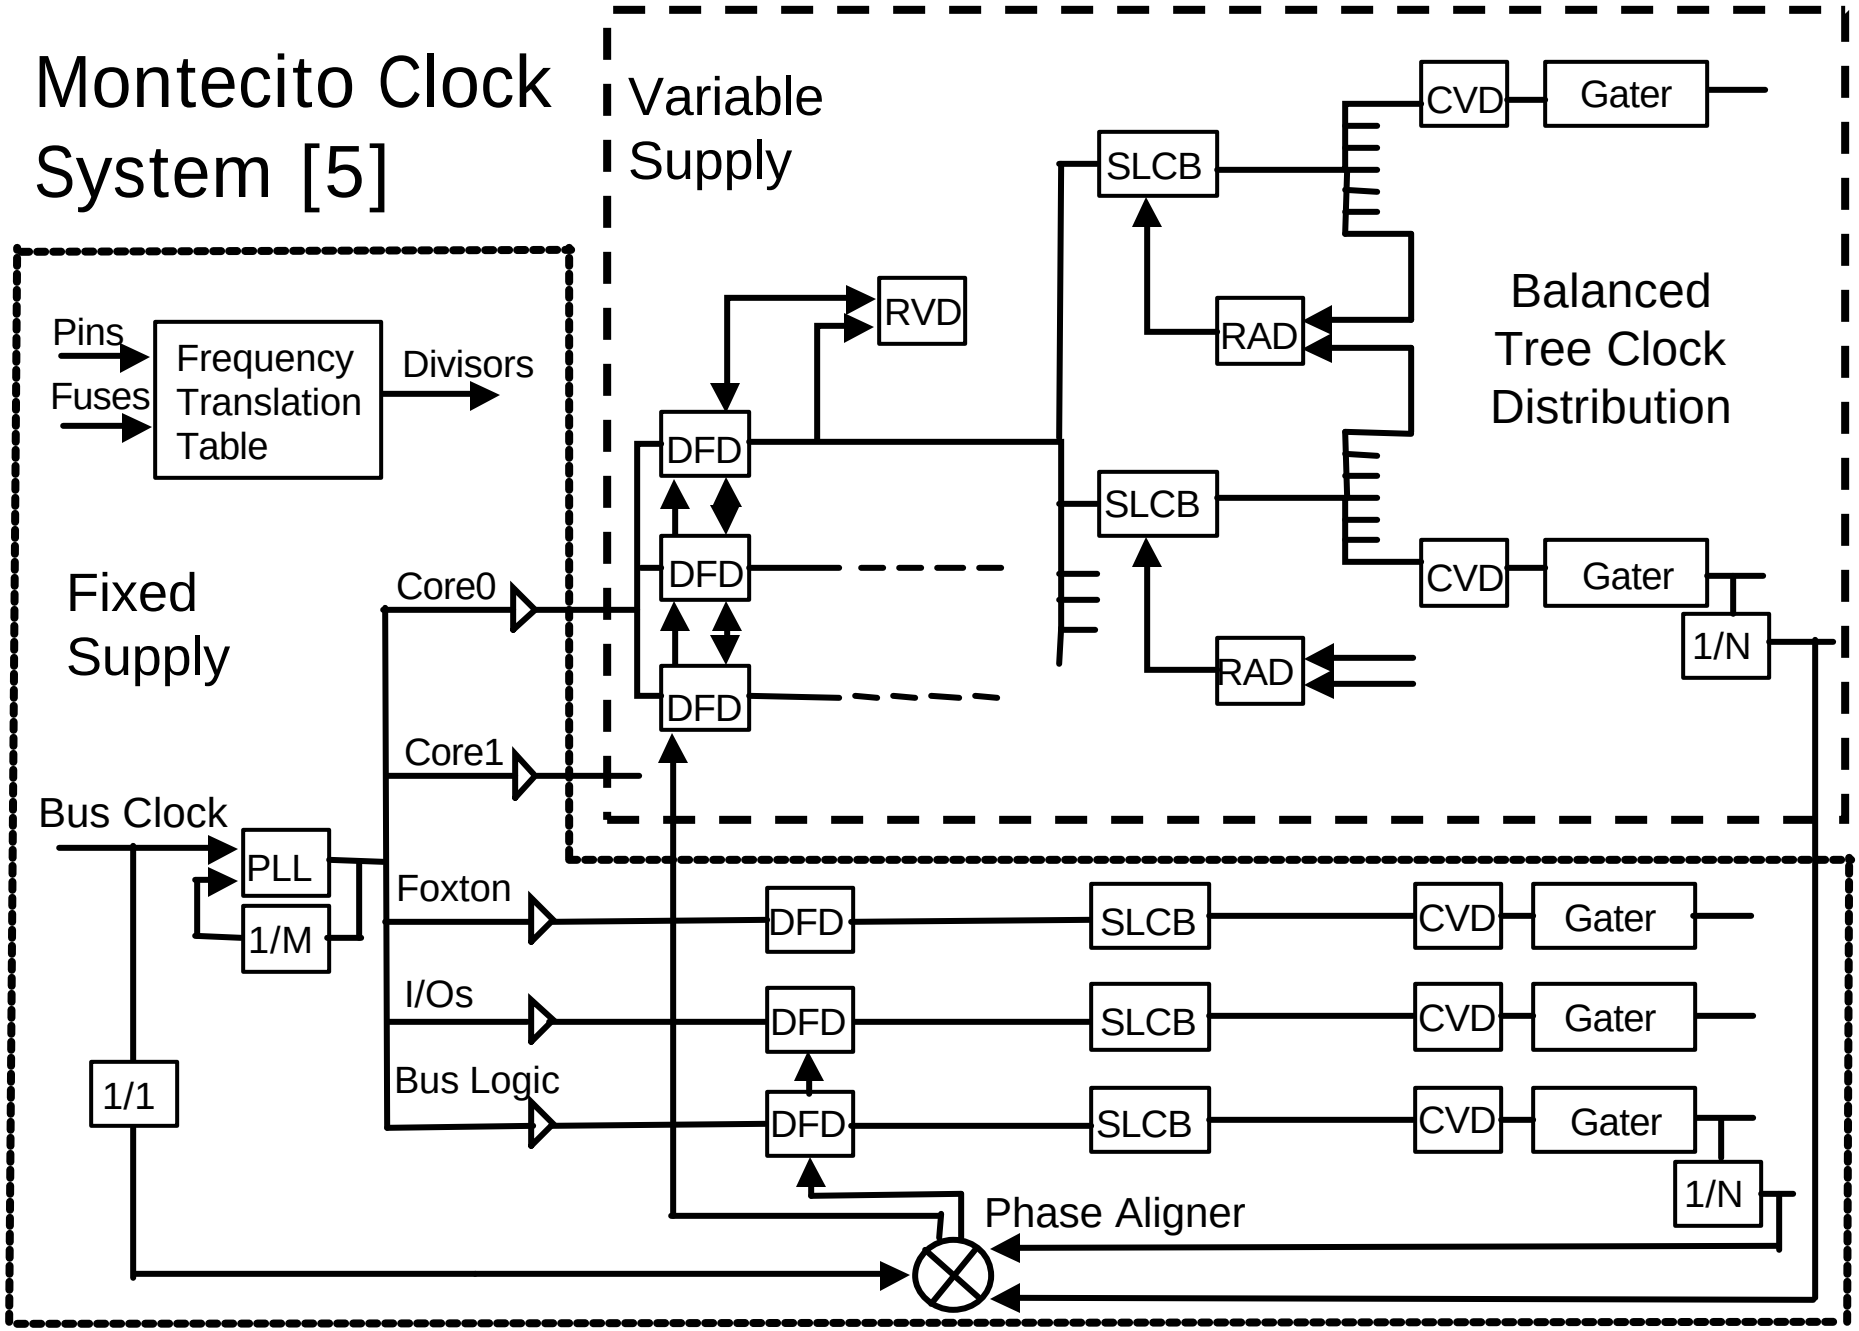

Montecito Clock Distribution [6]

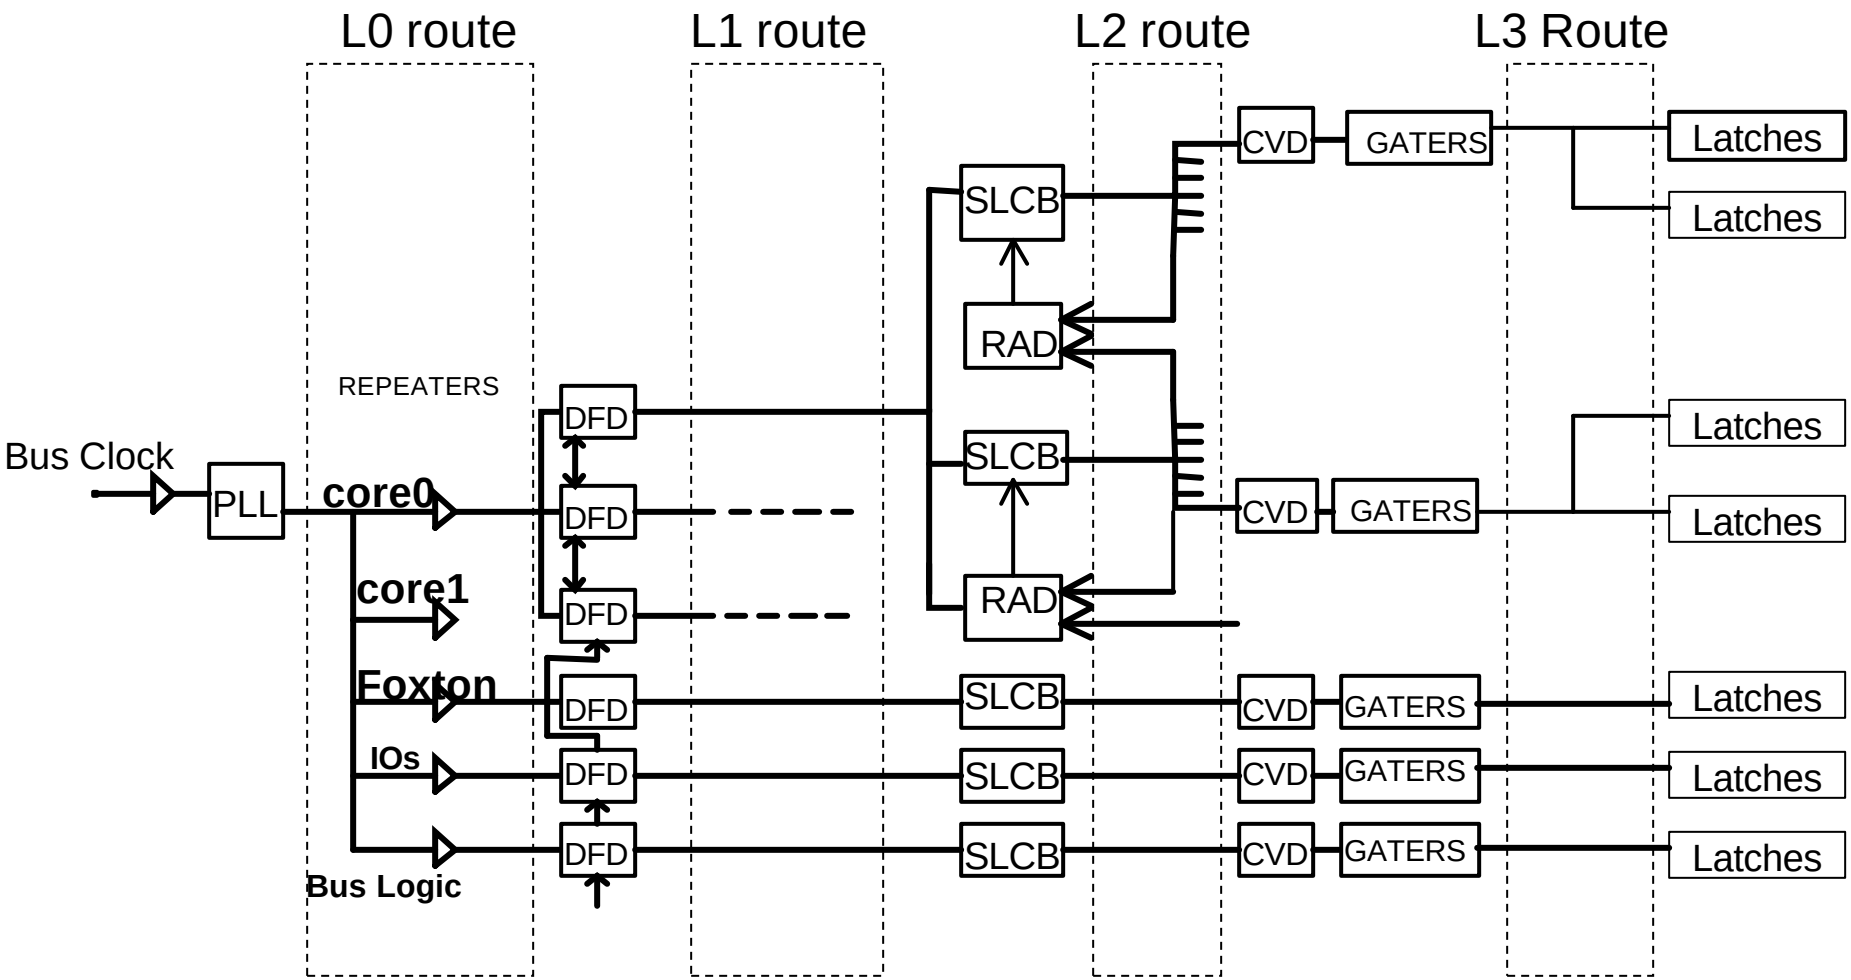

Single-Ended vs. Differential Clocks

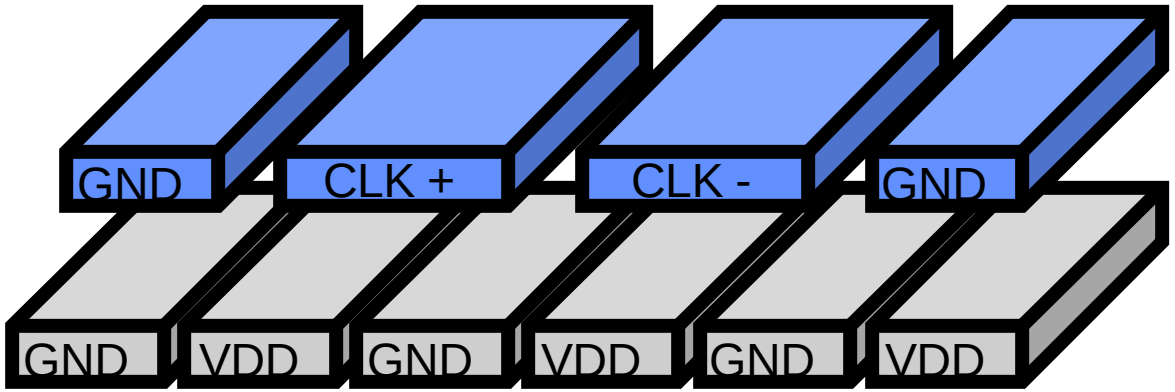

- Differential clock
  - Lower skew
  - High power
  - Longer distance between repeaters

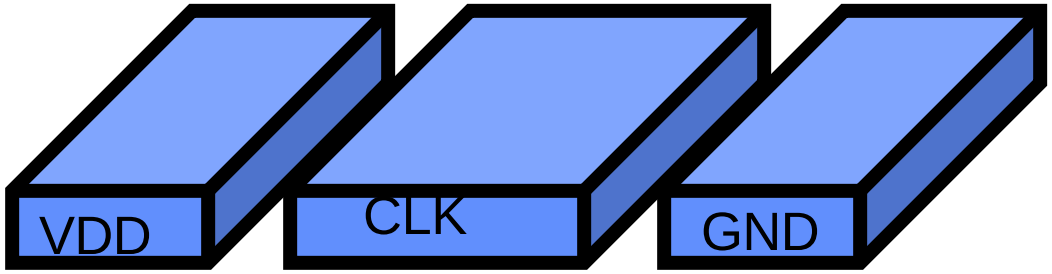

- Single-ended clock
  - Lower power
  - Need sharp edges to control skew

Dense vs. Sparse Grid Tiles

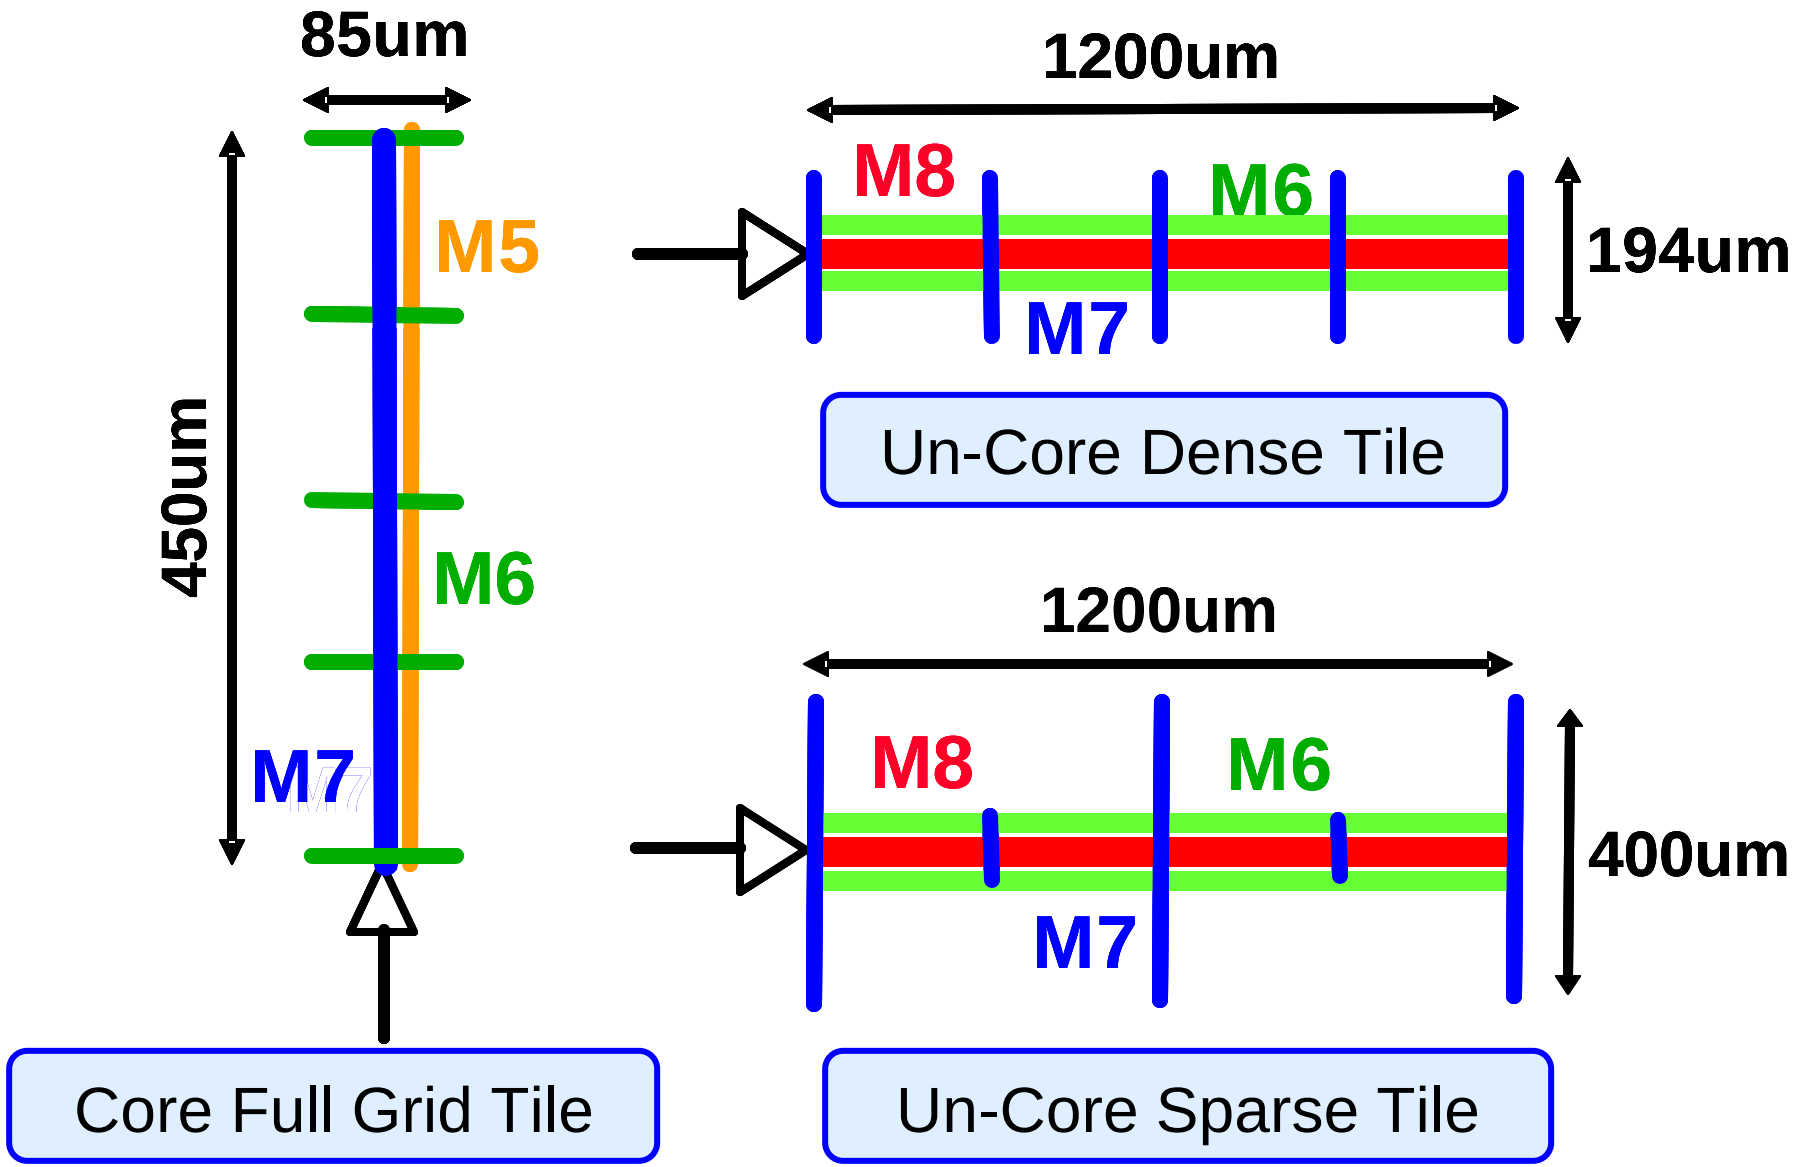

Xeon® Processor Package

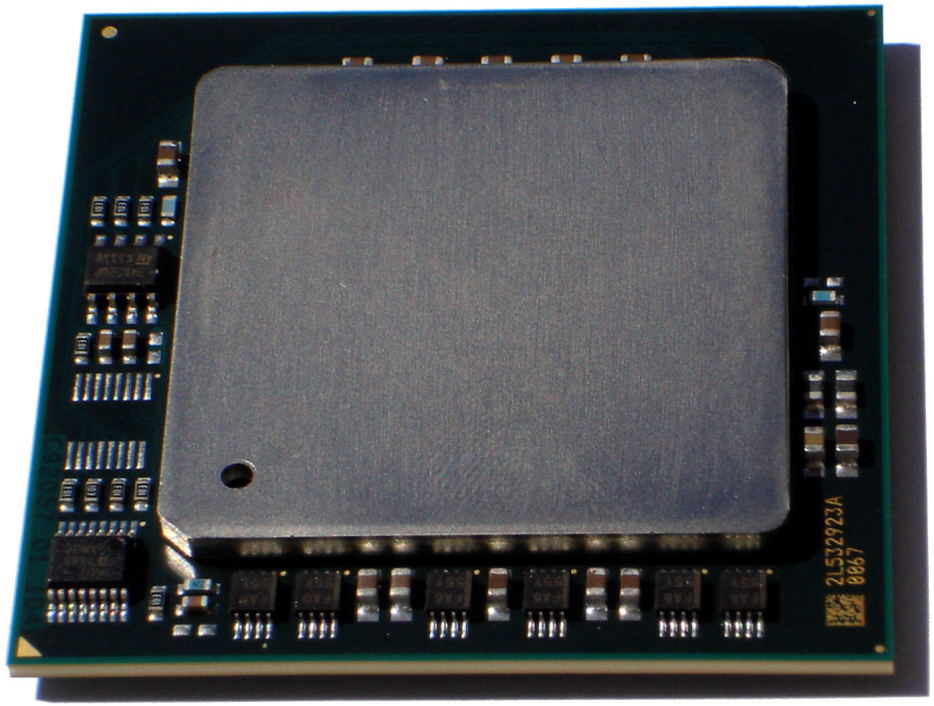

- 12 layers organic package (53.3 mm/side)
- 4-4-4 stacking
- Integrated heat spreader (38.5 mm/side)
- 604 total pins
- 366 signal I/Os
- System management components and decoupling capacitors on package

Itanium® 2 Processor Package

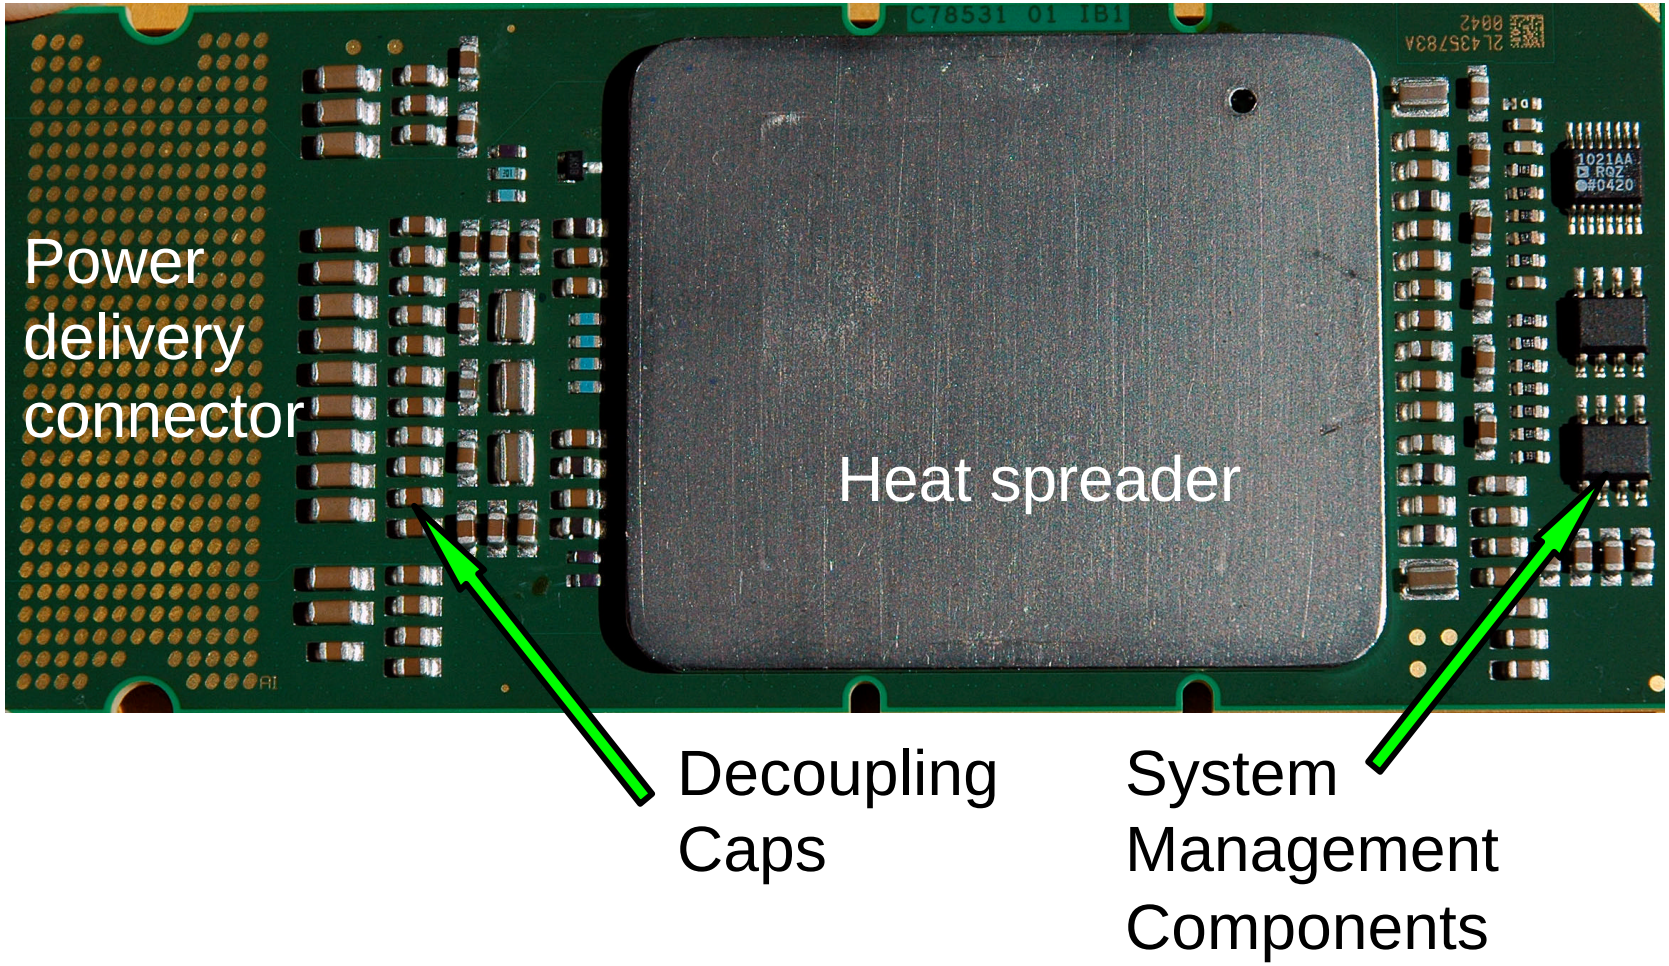

Intel® Core™ Duo Processor Package

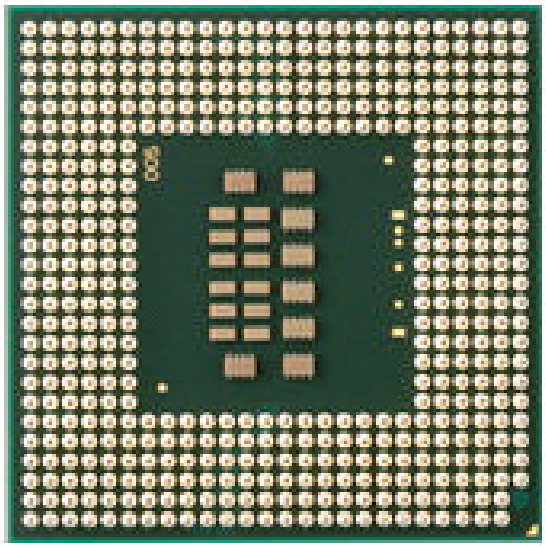

Design for Test and Debug Features

Die-level DFT/DFM

- Parallel structural core test with XOR
- Scan and observability registers (scan-out)
- Three TAP controllers (core0, core1, uncore)
- Within-die process monitors
- On-die clock shrink

L3 cache DFT/DFM

- Built-in pattern generator (PBIST)
- Programmable weak-write test
- Low-yield analysis
- Stability test mode
- 32-entry cache line disable (Pellston)

FSB DFT/DFM

- I/O loopback
- I/O test generator

Itanium® 2 Thermal Map

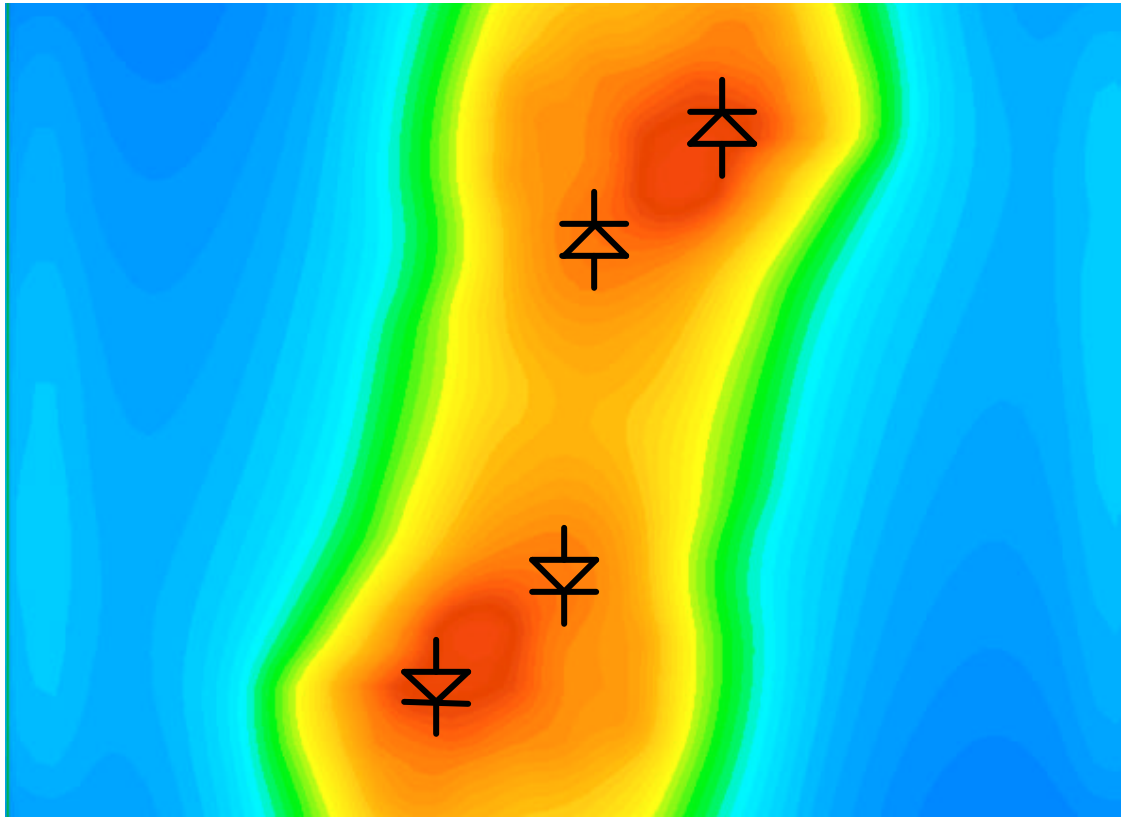

- Two thermal sensors per core
- Mux thermal diodes into VCOs to measure temp

Xeon® Infra-red Emission Image

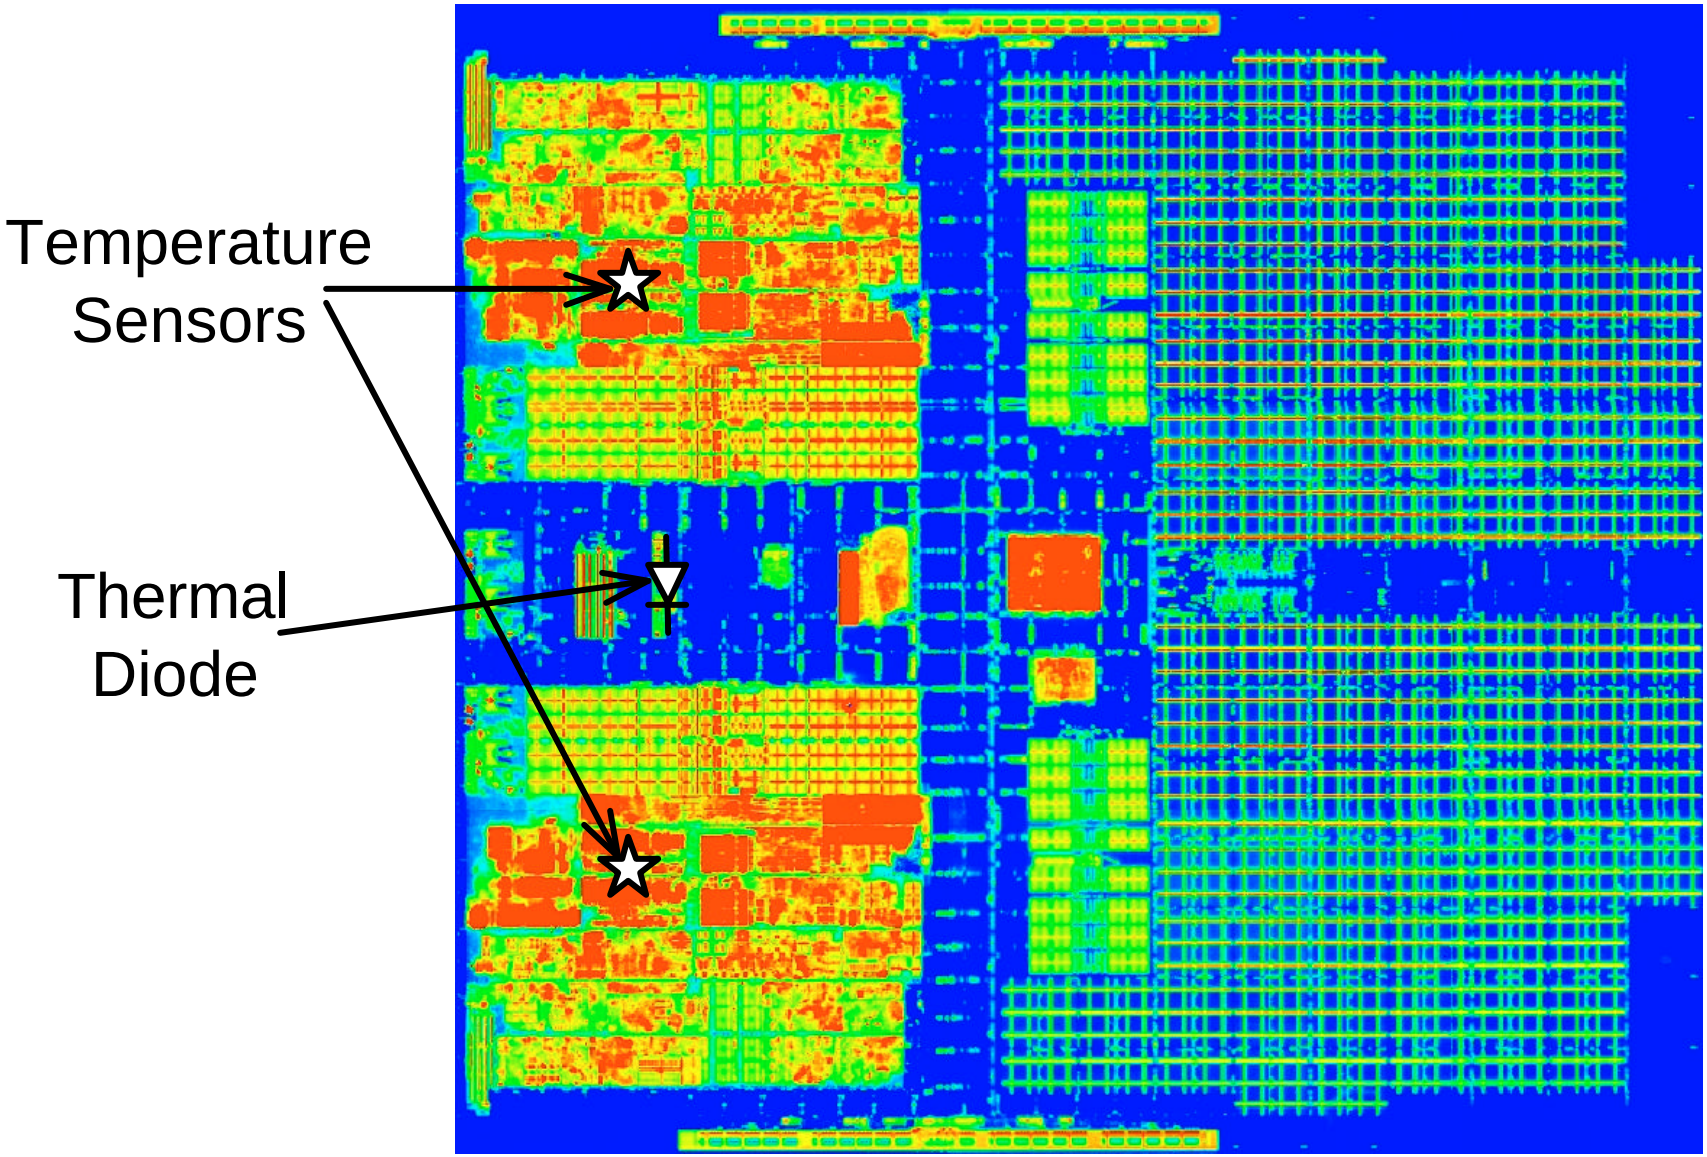

Potential Multi-Core Thermal Control

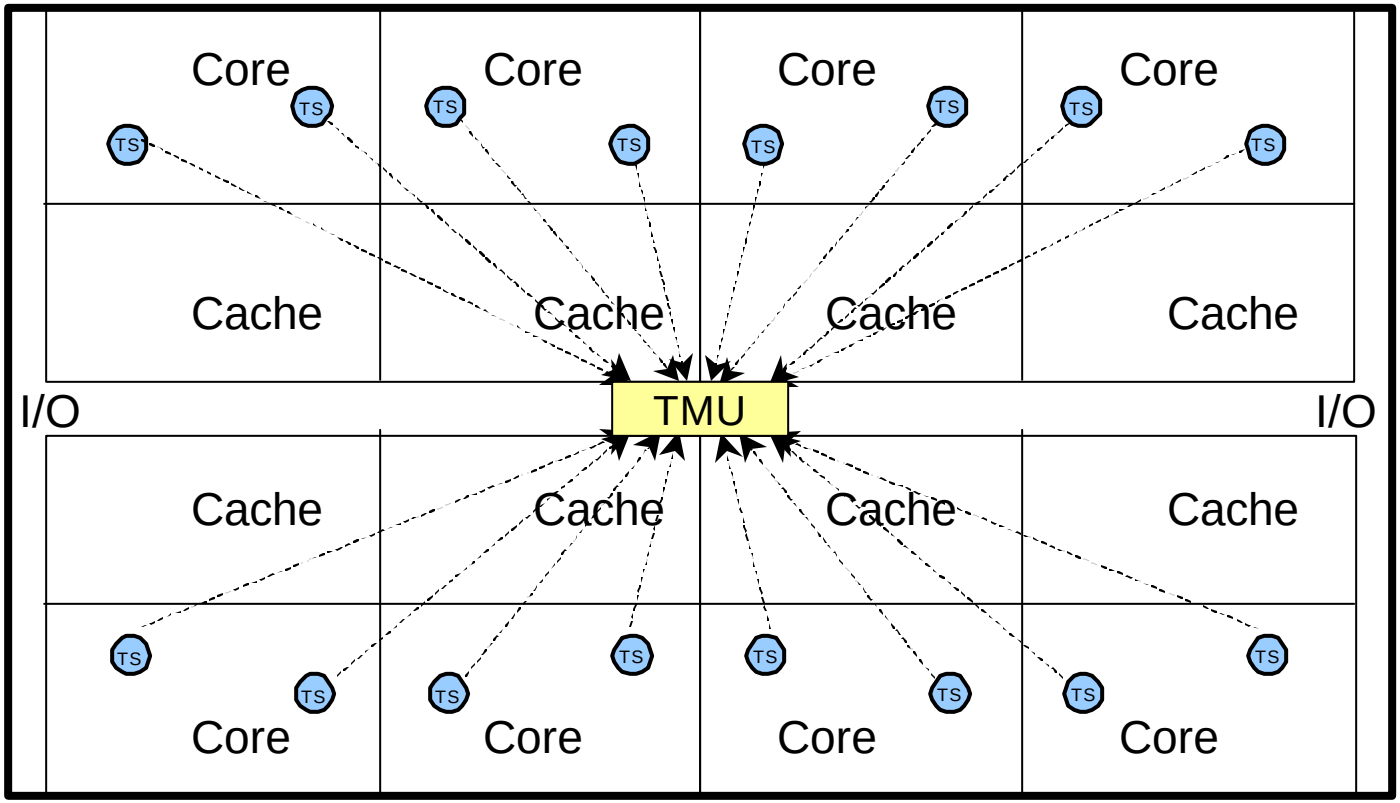

- Multiple core designs require a central thermal management unit and an efficient mechanism for transmitting thermal sensor measurements [8]

Summary

- Dual-core processors cover the entire compute spectrum, from laptops, to desktop and servers
- Moore's Law applied to multi-core processors:
  - Process technology scaling enables number of cores and cache size to double every two years
- Power and leakage reduction circuit techniques are essential for multi-core processors
  - Massive Long-Le usage **?** many slow transistors and a few fast devices
  - Cache sleep and shut-off modes (static and dynamic)
- Multiple on-die clock and voltage domains are required to control active power and leakage
  - Need new verification tools and automated checks

References

[1] S. Rusu, *et al*., “A Dual-Core Multi-Threaded Xeon® Processor with 16MB L3 Cache,” *ISSCC Dig. Tech. Papers*, Paper 5.3, Feb. 2006.

[2] S. Naffziger, *et al*., “The Implementation of a 2-Core, Multi-Threaded Itanium® Family Processor,” *IEEE J. Solid-State Circuits*, pp. 197-209, Jan. 2006.

[3] R. Korner, “Yonah and Sossaman Processor Briefing”, *Intel Developer Forum*, Fall 2005

[4] S. Tam, *et al*., “Clock Generation and Distribution of a Dual-Core Xeon® Processor with 16MB L3 Cache,” *ISSCC Dig. Tech. Papers*, Paper 21.2, 2006.

[5] T. Fischer, *et al*., “A 90-nm variable frequency clock system for a power managed Itanium® architecture processor,” *IEEE J. Solid-State Circuits*, pp. 217–227, Jan. 2006.

[6] P. Mahoney, *et al.,* “Clock Distribution on a Dual-core Multi-threaded Itanium™-Family Microprocessor,” *ISSCC Dig. Tech. Papers*, 2005

[7] S. Rusu, “Timing analysis of high-speed VLSI designs - Trends and challenges,” *International Workshop on Timing Issues (TAU),* 1997

[8] S. Rusu and S. Tam, “Apparatus for thermal management of multiple core microprocessors,” US patent 6,908,272, issued 7/21/2005


INSTITUTE OF ELECTRICAL AND ELECTRONICS ENGINEERS
ISSCC
UNIVERSITY OF PENNSYLVANIA